


PATENT
5681-71100
P8867

"EXPRESS MAIL" MAILING LABEL
NUMBER EV317117117US
DATE OF DEPOSIT

8-29-03
I HEREBY CERTIFY THAT THIS PAPER
OR FEE IS BEING DEPOSITED WITH
THE UNITED STATES POSTAL
SERVICE "EXPRESS MAIL POST
OFFICE TO ADDRESSEE" SERVICE
UNDER 37 C.F.R. §1.10 ON THE DATE
INDICATED ABOVE AND IS
ADDRESSED TO THE COMMISSIONER
FOR PATENTS, BOX PATENT
APPLICATION, P.O. BOX 1450,
ALEXANDRIA, VA 22313-1450.


Derrick Brown

SECURE TRANSFER OF HOST IDENTITIES

By:

James E. King and Martin P. Mayhead

B. Noel Kivlin
Meyertons, Hood, Kivlin, Kowert & Goetzl, P.C.
P.O. Box 398
Austin, TX 78767-0398

SECURE TRANSFER OF HOST IDENTITIES

BACKGROUND OF THE INVENTION

5 The present invention relates to the multiprocessor systems comprising a plurality of interconnected host systems.

10 In such a multiprocessor system, a host identity can be allocated to a processor domain of a host system, in order to control, for example, the distribution of software in accordance with a software licensing policy. A host identity can also be used to identify a host system to a network to which the host system is connected. It is important that a host identity is allocated to only one host system at a time in order, for example, to avoid the infringement of copyright in the software. For example, if the host identity were allocated to two host systems at a time, there might then exist
15 more copies of the software than are licensed to the user.

20 In this document, the term "host system" is used to define a processor domain, or a server, processor, or module having one processor domain, such that a host identity can only be active for one host system at a time.

25 In a multiprocessor system, where one or more of a set of available system resources can be allocated to perform a processing task, it may be desirable to reallocate the system resources to carry out the processing tasks. For example, if the system resource(s) originally allocated to perform the processing task fail(s), then it may be desirable to allocate one or more different system resources to enable further processing for that processing task. Reallocation of system resources in this manner can be described as "failover".

In view of the software licensing aspects identified above, it is important that an appropriate reallocation of the host identities takes place when system resources and/or processing tasks are reallocated in such a failover situation.

5 For example, if a server in a multiserver system were to fail, it may be desirable to allocate a "spare" server from a pool of one or more such spare servers. In such a case, the host identity allocated to the failed server would need to be transferred to the "spare" server. It would, therefore, be desirable to be able to reallocate the host identities between host systems.

10 Also, the transfer of host identities may be desirable to enable dynamic upgrading or modifying of a multi-user multiprocessor system. For instance, during use, an initially expected hardware resource requirement of a particular user may change. In such a situation, it would also be desirable to be able to reallocate host
15 identities between host systems so that users could be allocated hardware resources for addressing dynamically changing resource requirements.

 However, as indicated above, there is a need to ensure that a host identity is allocated to only one active host system at a time. The present invention seeks to
20 address the reallocation of host identities between host systems while ensuring that a host identity is allocated to only one active host system at a time.

SUMMARY OF THE INVENTION

Embodiments of the claimed invention can enable the secure transfer of host identities between host systems in a secure and reliable manner.

5

An embodiment of the invention can provide a method of transferring a first host identity between a first host system and a second host system. The first host system can be configured initially with the first host identity and the second host system can be configured initially with a second host identity, where a host identity
10 can belong to only one host system at a time. The method can include an administrator system, that is operable to hold an additional host identity, securely connecting to the first host system and passing the additional host identity to the first host system. The first host system can be operable to reconfigure itself to use the additional host identity and to pass the first host identity to the administrator system.
15 The administrator system can be operable to securely connect to the second host system and to pass the first host identity to the second host system. The second host system can be operable to reconfigure itself to use the first host identity and to pass the second host identity to the administrator system.

20 In this manner, the secure transfer of the first host identity from the first to the second host system can be achieved without a risk of the duplication of the host identity.

On return of the second host identity to the administrator system, this can be
25 placed in a pool of host identities for use as an additional host identity.

Alternatively, if required, the administrator system can be operable securely to connect to the first host system and to pass the second host identity to the first host system. The first host system can then be operable to reconfigure itself to use the

second host identity and to pass the additional host identity back to the administrator system. In this manner, the secure exchange of the first and second host identities between the first and second host systems can be achieved without a risk of the duplication of the host identities.

5

An embodiment of the present invention can also provide a computer system operable to transfer host identities between a first and second host system.

10 An embodiment of the invention can also be in the form of instructions for causing a host identity to be transferred between host systems, where a host identity can belong to only one host system at a time.

The instructions can be provided on a carrier medium. The carrier medium could be a storage medium, such as a solid state, magnetic, optical, magneto-optical or
15 other storage medium. The carrier medium could be a transmission medium such as broadcast, telephonic, computer network, wired, wireless, electrical, electromagnetic, optical or indeed any other transmission medium.

20 Further aspects and advantages of the invention will become apparent from the following description of particular embodiments.

BRIEF DESCRIPTION OF THE DRAWINGS

Embodiments of the present invention will be described hereinafter, by way of example only, with reference to the accompanying drawings in which like reference
5 signs relate to like elements and in which:

Figure 1 is a schematic representation of an architecture of a multiprocessor system for supporting a web site;

10 Figure 2 is a schematic representation of a racking system incorporating an example of a carrier in the form of a rack-mountable shelf according to a first example;

Figure 3 is a front view of an example of a carrier in the form of a shelf of
15 Figure 2;

Figure 4 is a rear view of an example of the shelf of Figure 2;

Figures 5A, 5B and 5C are schematic perspective views and a schematic
20 exploded view respectively of an example of an information processing cartridge for mounting in the shelf of Figure 2;

Figure 5D is a schematic perspective view of an example of an information processing cartridge for mounting in the shelf of Figure 2;
25

Figure 6 is a schematic perspective view of an example of a combined switch and service processor module for mounting in the shelf of Figure 2;

Figure 7 is a schematic perspective view of an example of a power supply module for mounting in the shelf of Figure 2;

Figures 8A, 8B and 8C are a schematic plan view and schematic perspective views, respectively, of an example of the chassis and midplane of the shelf of Figure 2;

Figure 8D is a schematic perspective view of a part of a base member of the shelf of Figure 2;

10

Figures 9A, 9B and 9C are schematic front, top and rear views, respectively, of an example of a midplane of the shelf of Figure 2;

Figure 10 is a schematic cross section view through the shelf of Figure 2;

15

Figure 11 is a functional block diagram of an example of an information processing subsystem for the information processing cartridge of Figure 5;

Figure 12 is a functional block diagram of an example of an information processing subsystem for the combined switch and service processor module of Figure 6;

Figure 13 is a functional block diagram of an example of a subsystem for the power supply unit of Figure 7;

25

Figure 14 is a functional block diagram showing the connectivity between the components of the shelf of Figure 2;

Figure 15 is a functional block diagram showing the external connectivity of the shelf of Figure 2;

Figure 16 is a schematic representation of a shelf showing the external
5 connections from the shelf of Figure 2;

Figure 17 is a schematic representation of a rack mounted system comprising a plurality of such shelves;

10 Figure 18 is a functional block diagram of a further example of an information processing subsystem for the information processing cartridge of Figure 5;

Figure 19 illustrates part of a multi-domain computer system;

15 Figure 20 is a schematic representation of components of the system of Figure 17 for explaining specific examples of the present invention;

Figure 21 illustrates one method of operation in accordance with an embodiment of the invention; and

20

Figure 22 illustrates an example of operation in accordance with the present invention; and

Figure 23 illustrates another example of operation in accordance with the
25 present invention.

While the invention is susceptible to various modifications and alternative forms, specific embodiments are shown by way of example in the drawings and are herein described in detail. It should be understood, however, that drawings and

detailed description thereto are not intended to limit the invention to the particular form disclosed, but on the contrary, the invention is to cover all modifications, equivalents and alternatives falling within the spirit and scope of the present invention as defined by the appended claims.

5

DESCRIPTION OF PARTICULAR EMBODIMENTS

Embodiments and examples are described hereafter by way of example with reference to the accompanying drawings.

5

Shown in Figure 1 is an example of an application of a high capacity multiserver system 1 for implementing a network-connected web site such as, for example, an airline reservation system on the World Wide Web.

10 As shown in Figure 1, an external network 3 (e.g., the Internet) for communicating with a user 5 can be connected to gateways 7 which can be connected to an entry edge server group 9 implemented by a web farm. The entry edge server group 9 forms an interface to the external network 3. The entry edge server group 9 can then be connected by switches 11 and a firewall 13 to a web edge server group 15
15 that can also be implemented as a web farm as shown in Figure 1. The web edge server group 15 can serve to cache web pages that are readily accessible to users 5 accessing the system 1 from the external network 3, for example for checking flight times, etc. The web edge server group can comprise a number of blade server (BS) shelves and a number of network addressable storage (NAS) shelves for storing
20 critical data. The web edge server group 15 can be further connected by a further firewall 17 to a plurality of application servers 19, which can be responsible for, for example, processing flight reservations. The application servers 19 can then be connected via a further firewall 21 to computer systems 23, 25, for example, e-commerce services including financial services for receiving and processing payment
25 for airline reservations.

As will be appreciated, the server system described above with reference to Figure 1 is only an example of a possible application for a multiprocessor server system. Multiprocessor server systems have many different applications and the

present system is not limited to being applicable for use in only one or a limited number of such applications, rather multiprocessor server systems as described herein are operable for use in many different applications. A non-exhaustive list of such alternative applications includes: e-commerce web server systems; 5 telecommunications network server systems; LAN application and file server systems and remote vehicle control systems.

With reference to Figure 2, there is shown a schematic perspective representation of a rack system 31 as viewed from the front including left and right 10 front uprights 32 and 33 and left and right rear uprights 34 and 35. The uprights can be formed with apertures for receiving shelf fixings (e.g., screws, bolts, clips, etc., for mounting brackets, slides, rails, etc.).

Also shown in Figure 2 is an example of a blade server shelf 41 mounted in 15 the rack system 31. The shelf 41 forms a carrier configured to carry a plurality of information processing cartridges 43 located side by side along the shelf.

The term "shelf" is used herein in a conventional way to describe a structure that is mountable in rack system 31 and is configured to carry one or more 20 components to form at least a part of a rack-mountable system. In the present example, the shelf 41 is three-dimensional, having a height (H), width (W) and depth (D). In the present example, one dimension (hereinafter described as the height, H) is smaller than the other dimensions (hereinafter described as the depth, D, and the width, W) to facilitate mounting of the shelf within the rack system 31. It will be 25 appreciated that although the width and depth are typically constrained by the dimensions of the racking system for which the shelf is designed, there is more freedom as regard the height, subject to taking account of any appropriate standards and packaging considerations.

Each of the information processing cartridges contains at least one processor. Each information processing cartridge in the present example is operable as a server. In the described examples, the information processing cartridges are configured as robust enclosed modules.

5

In the example to be described in more detail presently, the information processing cartridges, when aligned in the carrier shelf, look like rectangular slabs, or blades. Accordingly, an information processing cartridge can be described as a blade.

The information processing cartridges 43 comprise information processing modules enclosed in an enclosure, or housing, so that the information processing modules have the form of cartridges. Also, as the information processing cartridges are to operate as computer servers in the example described in more detail presently, an information processing cartridge 43 can also be described as a server blade. Accordingly, in the context of this example, the terms module, cartridge and blade are used interchangeably.

The illustrated example of a shelf 41 is configured to carry sixteen information processing cartridges 43, each of which is removably mountable in a respective opening 45 in the front of the shelf, whereby the information processing cartridges can be inserted into and removed from the front of the shelf 41 without removing the shelf 41 from the rack system 31.

In the present example, the shelf 41 comprises a three-dimensional, generally rectangular, enclosure, or housing, 47 that is suitable for mounting in generic racking systems including both 4-post and 2-post systems. It can be mounted on fixed rigid rack mounting ears and/or a simple slide/support system. The present example is designed for standard 19"-wide racking (1"=25.4mm) as defined, for example, in the well-known IEC297 and EIA 310 specification standards with height corresponding to the so-called 3U (3 standard unit) height. For mounting such a 3U unit in such a 19"-

wide racking system, with a depth of, say 25" or 30", the enclosure can be arranged with a height of up to about 130.5mm, a width of up to about 445 mm and a depth, including all hardware and fascias, but excluding cable management, of up to about 635mm, with the depth from the front-most point of a fascia to a rear I/O connector panel of a rear mounted Field Replaceable Unit (FRU) of about 610mm. Of course, other examples designed for other racking systems could have different dimensions.

This example of a shelf 41 has a single enclosure, or housing, 47 that houses a number of modular units or subsystems, the majority of which are replaceable in the field and are therefore known as Field Replaceable Units (FRUs). These modular units include the information processing cartridges 43.

The shelf enclosure 47 can be fabricated from sheet material (e.g., from steel sheet) to form a chassis portion 49 that includes a base 51, two sides 53 and 55, a front 57 and a rear 59. The word "front" as used here is merely used as a label herein to refer to the face, or wall 57 of the enclosure that is located at the main access side of the rack system 31 in use when the shelf is mounted therein. Similarly, the words "rear" and "side" are merely used as labels herein to refer to the faces, or walls 59, 53 and 55 that, in use, are located at those respective positions when the shelf is mounted in the rack system 31.

The openings 45 can be formed in the front face 57 for receiving the information processing cartridges 43 and, as will be explained later, apertures can also be formed in the rear face 59 for receiving further FRUs. The enclosure can further include a removable top cover 61 that can be secured to the chassis portion 49 by suitable fastening (e.g., screws). The apertures in the front and rear faces 57 and 59 allow at least some of the FRUs to be inserted into and/or removed from the shelf enclosure 47 via the front or the rear thereof, as appropriate, without removing the shelf from the racking. Access to components mounted in the shelf that are not

accessible via one of the apertures in the front 47 and rear 59 faces can be achieved by removing the shelf enclosure 47 from the racking system 31 and then removing the top cover 61 of the shelf enclosure 47.

5 Figure 3 is a front view of an example of a shelf 41 for a first example. A plastic front bezel 63 can be provided that fits on the front face 57 (shown in Figure 2) of the chassis 49 of the shelf enclosure 47. The front bezel 63 can be formed as a unitary removable part that spans the whole width and height of the front of the shelf enclosure 47. The front bezel 63 could alternatively include a number of separate
10 components, or moldings. The front bezel can include a peripheral portion 64 that can provide areas for corporate and product branding marks, for identification and numbering for the information processing cartridge and for a bar code label (all not shown). One or more apertures 65 can be formed in the peripheral portion 64 of the bezel 63. The apertures 65 in the bezel can be arranged to align with one or more
15 apertures (e.g. a slot (not shown in Figure 3) in the front face of the chassis. In use, air can pass through the apertures 65 to flow into the shelf enclosure 47 to reach FRUs that are mounted in the shelf enclosure 47 through the rear face 59 thereof. Air flowing through the aperture 65 flows into a plenum chamber 66 (not shown in Figure 3) to flow past the processing cartridges 43 to reach rear mounted FRUs. A central
20 area 67 of the front bezel 63 can be open allowing access to the openings 45 in the front face 57 of the shelf enclosure 47 for insertion and removal of the information processing cartridges 43. Where no active module is mounted in a location for an information processing module, a blanking panel, or filler panel, such as the blanking panel 44, can be located in that location. LED indicators 69 can be mounted on a
25 system indicator printed circuit board (not shown) behind a designated area of the bezel to provide an indication of system status via light guides incorporated into the bezel. A further system indicator board (also not shown) carrying LED indicators can be provided inside the shelf enclosure to be visible from the rear thereof.

As mentioned above, in the present example of a shelf, up to sixteen information processing cartridges 43 can be installed in respective openings 45 in the front face 57 thereof. The number of information processing cartridges 43 actually installed in any installation is dependent upon the system configuration required.

5 Various features relating to the information processing cartridges 43 that are shown in Figure 3 will be described later.

Figure 4 illustrates the rear of the shelf unit of Figures 2 and 3. This shows two different types of FRU 71 and 81 (4 units in total) that have been inserted into

10 respective apertures 72 and 82 in the rear of the shelf enclosure 47. The FRUs shown in Figure 4 include two Combined Switch and Service Processors (CSSPs) 71 and two Power Supply Units (PSUs) 81. Various features shown in Figure 4 will be described later.

15 Before proceeding with a more detailed description of each of the FRUs 43, 71 and 81 introduced so far and of the construction of the shelf 41, there follows a brief description of an information processing cartridge 43, a CSSP 71 and a PSU 81 with reference to Figures 3, 4, 5, 6 and 7.

20 Figure 5A provides a perspective view, partly from the rear, of an information processing cartridge 43. Figure 5B provides a perspective view, partly from the front of the same information processing cartridge 43. Figure 5C provides an exploded perspective view of the construction of the information processing cartridge 43. Here is it to be noted that the term "rear" is applied in the context of the position, when

25 installed, of the information processing cartridge, with respect to the shelf 41 (i.e. in this case the "rear" of the information processing cartridge 43 is the innermost part of the information processing cartridge when it is inserted in the shelf 41). Likewise, "front" refers in the present context to the outermost part of the information processing cartridge when it is inserted in the shelf 41.

With reference to Figures 3, 5A and 5B, it will be noted that information processing cartridges are three-dimensional, having a height (h), width (w) and depth (d). If, as in the present example, the information processing cartridges are to be arranged in a one-dimensional array (a row) across the shelf, then efficient packing for the information processing cartridges is achieved where one dimension (here the width, w) is smaller than the other dimensions (here the depth, d, and the height, h). In a particular example, the enclosure of a processing cartridge 43 has a height h, width w and depth d of 115mm, 26mm and 315mm, respectively, although of course, other examples may have different dimensions.

It will be noted that an enclosure 101 of the present example of an information processing cartridge 43 has six, generally rectangular, faces. For convenience only, the face that is visible from the front of the racking when an information processing cartridge 43 is mounted in the shelf 41 is known as the front face 102. The opposite face is known as the rear face 103. In the present example these two faces, as well as top and bottom faces 104 and 105, and the side faces 106 and 107 have the shape of elongate rectangles.

Although in the present example the information processing cartridges have six generally rectangular faces, it will be appreciated that other examples could have other configurations. For example, rather than having a generally rectangular shape, the side faces of another example of an information processing cartridge could have the general shape of a triangle (whereby the information processing cartridge may then only have five faces), a pentagon (whereby the information processing cartridge may then have seven faces), and so on. Indeed, one or more or all of the edges could be curved. However, it will be appreciated that the present configuration provides advantages for example, in terms of manufacturing, engineering and packing density within a shelf 41.

In this example, the information processing cartridge enclosure 101 is fabricated from two housing portions. The first housing portion 109 can be fabricated from, for example, a metal (e.g., pressed steel) and can include one side face 106, the rear face 103 and part of the top face 104. The part of the top face formed by the metal portion 109 is given the reference 1041 in Figure 5A. The second housing portion 110 can be fabricated from, for example, a plastics material and can include the other side face 107 and the front faces 102 and bottom face 105 and the remaining part 1042 of the top face 104. In the present example, a plastics material used is PolyCarbonate Acrylonitrile Butadiene Styrene (PCABS), however many other plastics materials such as other ABSs and nylons may be used. The metal portion 109 also includes a part 1021 corresponding to the front face 102 and a part 1051 corresponding to a part of the bottom face 105. The parts are covered by the plastics portion 110 when enclosure 101 is fully assembled. These parts are visible in the exploded view of Figure 5C. The two housing portions 109 and 110 are secured to one another by fixings such as screws 118.

As shown in Figure 5A, a groove 108 runs along the length of the plastic top face portion 1042. This groove 108 is provided for interfacing with a guide member of the shelf chassis (not shown in Figure 5A, but see Figures 8B and 8C). A similar groove 113 is formed in the bottom face 105 of the plastics portion 110 as shown in Figures 5B and 5C. This groove 113 is provided for interfacing with a guide member of the shelf chassis (not shown in Figures 5B and 5C, but see Figure 8D).

The provision of the enclosure 101 means that the information processing cartridge 43 can safely be handled by an operator who is not a skilled technician. Also, through the use of the enclosure 101, the information processing cartridge is a robust unit that protects its inner workings from the outside environment and vice versa. The use of a conductive enclosure, e.g., a metal enclosure, can result in the

information processing cartridge including its own electromagnetic shielding. To this end, the enclosure 101 can also be provided with an internal electromagnetic interference (EMI) shielding member 116 as shown in Figure 5C. The EMI shielding member can be located between the internal components 112 of the information processing cartridge 43 (not described in detail with reference to Figures 5A to 5C, but see Figures 11 and 18 below) and the plastics portion 110. The EMI shielding member can be secured to the plastics portion 110, for example by heat bonding or an adhesive. In other examples, the plastics portion could have a conductive layer deposited thereon, or conductive plastics material could be used. In this example, the EMI shielding member 116 is provided with electromagnetic interference (EMI) fingers 114 to ensure good contact with the shelf chassis and the adjacent components. These fingers 114 extend through EMI finger apertures 1141 in the plastics portion 110.

As shown in Figures 3, 5B and 5C, the information processing cartridge 43 incorporates an injector/ejector handle 111 on the front face 102 for facilitating insertion and latching of the information processing cartridge 43 within an aperture in the shelf 41.

As shown in those Figures, the handle 111 of the injector/ejector lever extends substantially the whole height of the front face of the information processing cartridge 43, thereby increasing the mechanical advantage and facilitating injection and ejection of the information processing cartridge 43. As is further shown in those Figures, the front face 102 of the information processing cartridge 43 has perforations 115, in the present example slits, to allow for airflow into the information processing cartridge 43. The front face part 1021 of the metal portion 109 has perforations corresponding to those in the front face 102 such that airflow into the information processing cartridge 43 is not impeded. It will be noted in Figures 3, 5B and 5C that the handle 111 is narrower in its middle than at its ends. This reduces any masking effect of the

handle 111 on the airflow to the perforations 115 in the front face of the information processing cartridge 43 and facilitates grasping of handle 111. The handle 111 can be bowed out from the front face of the information processing cartridge to further facilitate grasping thereof and to reduce any masking effect with regard to the 5 perforations. In other examples, the handle 111 could have an open frame-like structure to further facilitate airflow.

As shown in Figure 5A, the rear face 103 of the information processing cartridge 43 also has perforations 117 to allow for air to be exhausted from the rear of 10 the information processing cartridge 43. A fan can be located within the enclosure 101 of an information processing cartridge 43. In the present example the fan is a combined radial fan (blower) and heat sink to direct cooling air onto a processor of the information processing cartridge, which fan also encourages air movement within the enclosure 101. LED indicators 119 (see Figure 3 and 5B) can be provided on the front 15 face 102 of an information processing cartridge to indicate whether power is on, whether service intervention is required and whether the information processing cartridge 43 can be removed. A connector 120, for example a 40 way single connector attachment (SCA-2) connector (a small computer systems interface (SCSI) connector), can be provided at the rear of the information processing cartridge 43 for 20 electrical connection of the information processing cartridge 43 within the shelf 41. The connector 120 is advantageously able to withstand repeated removals and insertions of the enclosure 101 from and into a shelf 41. The connector arrangement can include a guide pin arrangement to prevent module misalignment during insertion of the information processing cartridge into the receiving location.

25

Thus an example of features and the construction of an information processing cartridge enclosure has been described. Although particular materials and constructions have been described, it will be appreciated that other examples could be employed. Indeed, it will be appreciated that this particular example relates to only a

possible form for the processing cartridge 43. An example of an alternative construction will now be described with reference to Figure 5D.

Figure 5D provides a perspective view, partly from the rear, of an alternative
5 information processing cartridge 43. Here it is to be noted that the term "rear" is applied in the context of the position, when installed, of the information processing cartridge, with respect to the shelf 41 (i.e. in this case the "rear" of the information processing cartridge 43 is the innermost part of the information processing cartridge when it is inserted in the shelf 41).

10

In this example, the information processing cartridge enclosure 101 is fabricated from pressed steel to form two chassis portions. The first portion 234 includes one side face 107, and part of each of the front and rear faces 102 and 103 and the top and bottom faces 104 and 105. The second portion 235 includes the other
15 side face 108 and the remaining part of each of the front and rear faces 102 and 103 and the top and bottom faces 104 and 105. The two chassis portions 234 and 235 meet at a groove 236 and are secured to one another by fixings (e.g., one or more screws, not shown). Grooves 236 run along the top and bottom faces 104 and 105 of the enclosure 101 and are provided for interfacing with guide rails of the shelf chassis
20 49 (not shown in Figure 5, but see Figure 8C). A cover portion that is secured to the chassis portion forms the other side face 106. It will be appreciated however, that in another example, the chassis portions could be joined at a position other than the groove 108, with the groove 108 being formed entirely in one of the chassis portions. Alternatively, the enclosure 101 may be constructed from a number of sheets of steel,
25 with each sheet forming one of the faces.

Figure 6 provides a perspective view, partly from the front, of a Combined Switch and Service Processor (CSSP) cartridge (or CSSP) 71. Here it is to be noted that the term "front" is applied in the context of the position, when installed, of the

CSSP cartridge 71, with respect to the shelf 41 (i.e. in this case the “front” of the CSSP cartridge 71 is the innermost part of the CSSP cartridge 71 when it is inserted in the shelf 41).

5 With reference to Figures 4 and 6, it will be noted that a CSSP cartridge 71 is three-dimensional, having a height (h), width (w) and depth (d). In a particular example, the enclosure of a CSSP 71 has a height h, width w and depth d of 43mm, 202mm and 278mm, respectively, although of course, other examples may have different dimensions.

10

 An enclosure 121 of present example of a CSSP cartridge 71 has six, generally rectangular, faces. For convenience only, the face that is visible from the rear of the racking when a CSSP cartridge 71 is mounted in the shelf 41 is known as the rear face 122. The opposite face is known as the front face 123. In the present example these
15 two faces, as well as side faces 126 and 127 have the shape of elongate rectangles. The top and bottom faces 124 and 125 are also rectangular, but not elongate in the manner of the front, rear, top and bottom faces. Although in the present example the CSSP cartridges have six generally rectangular faces, as for the information processing cartridges 43 it will be appreciated that other examples could have other
20 configurations.

 In this example, the CSSP enclosure 121 is fabricated from steel sheet to form a chassis portion that includes the bottom face 125, the front and rear faces 122 and 123 and the side faces 126 and 127. A cover portion that is secured to the chassis
25 portion forms the other top face 124. The cover portion is secured to the chassis portion by suitable fixings, for example one or more screws 128. It will be appreciated however, that in another example, other faces, or portions, of the enclosure could form the chassis and the cover portions. The provision of the enclosure 121 means that the CSSP cartridge 71 can safely be handled by an operator

who is not a skilled technician. Also, through the use of the enclosure 121, the switch cartridge is a robust unit that protects its inner workings from the outside environment and vice versa. The use of a conductive enclosure, e.g., a metal enclosure, means that the CSSP cartridge includes its own electromagnetic shielding. To this end the CSSP enclosure 121 is provided with EMI fingers 129 to ensure good contact with the shelf chassis and the adjacent components.

As shown in Figure 4, the CSSP cartridge 71 incorporates two D-shaped handles 132 to facilitate insertion and removal of the CSSP cartridge 71 with respect to an aperture 72 in the rear face of the shelf enclosure. A latch member 131 can be pivotably mounted on a plate that can be secured (e.g., using screws) to the rear face of the shelf enclosure. The latch member 131 is configured to engage one of the handles 132 and to secure the CSSP cartridge 71 in place. In other examples, the CSSP cartridge 71 could be provided with an injector/ejector handle in a manner similar to the information processing cartridge. As shown in Figure 6, the front face 123 of the CSSP cartridge 71 has perforations 133 to allow for airflow into the CSSP cartridge 71. As shown in Figure 4, the rear face 122 of the CSSP cartridge 71 has perforations 135 to allow for air to be exhausted from the rear of the CSSP cartridge 71.

At least one fan can be located, for example behind the perforated portion 135 of the rear face, in a CSSP cartridge 71 to channel cooling air through the CSSP cartridge 71 from the front to the rear. In this particular example shown, two fans are provided, one behind each set of perforations 135. LED indicators 137, as shown in Figure 4, can be provided on the rear face 122 of the CSSP enclosure 121 to indicate whether power is on, whether service intervention is required and whether the switch can be removed. Additional link status indicators can be provided integral to 2x4 stacked RJ-45 connectors 139, also shown in Figure 4. As shown in Figure 4, electrical connections 141 can be provided at the front face of the CSSP (i.e. on the

face that in use is inside the shelf enclosure 47). Suitable connections for use in the present example include a connector for power connections, a connector for serial management data connections and a connector for information connections. In the present example, information connections are implemented using an Ethernet
5 information communication protocol, e.g. at 1 Gigabit (Gb). However other protocols could equally be used, for example the Infiniband information communication protocol. The connector arrangement can include a guide pin arrangement to prevent module misalignment during insertion of the CSSP module into the receiving location. For this purpose, guide pin holes 142 can be provided on the front face 123
10 into which guide pins may pass to aid module alignment.

In the present example, up to two CSSPs 71 can be mounted at any one time at the rear of the shelf unit in corresponding apertures 72 in the rear face of the shelf enclosure 47. The number of CSSPs 71 provided in any particular implementation
15 depends upon system configuration and the need, or otherwise, for redundancy.

It will be appreciated that one possible construction of the CSSP cartridge 71 has been described and that as for the information processing cartridge 43, other examples could employ other materials and/or constructions.
20

Figure 7 provides a perspective view, partly from the front, of a power supply unit (PSU) cartridge 81. Here it is to be noted that the term "front" is applied in the context of the position, when installed, of the PSU cartridge 81, with respect to the shelf 41 (i.e. in this case the "front" of the PSU cartridge 81 is the innermost part of
25 the PSU cartridge 81 when it is inserted in the shelf 41).

With reference to Figures 4 and 7, it will be noted that PSU cartridge 81 is three-dimensional, having a height (h), width (w) and depth (d). In this particular example, the order to provide for a dense packing of the FRUs in the rear of the shelf

41, the PSU cartridge 81 has two dimensions (hereinafter described as the width, w, and the depth, d) that are generally similar. In a particular example, the enclosure of a PSU cartridge 81 has a height h, width w and depth d of 83mm, 202mm and 276mm, respectively, although of course, other examples may have different dimensions.

5

An enclosure 145 of present example of a PSU cartridge 81 is of generally oblong shape, but has the "top" "front" edge cut away to form an additional "top" "front" sloping face. The enclosure 145 therefore has five, generally rectangular, faces and two faces of generally rectangular shape with one corner cut away. For convenience only, the face that is visible from the rear of the racking when the PSU cartridge 81 is mounted in the shelf 41 is known as the rear face 146. The opposite face is known as the front face 147. In the present example these two faces and the two side faces 150, 151 are of elongate, generally rectangular shape with one corner cut away, given that the width and depth of the PSU cartridge are similar, whereas the top and bottom faces 148, 149, although still rectangular, are not, in this example, notably elongate. A top front face 148a is present at the top front of the enclosure. Thus the front of the enclosure is sloped at the top edge. As for the information processing cartridges 43, however, it will be appreciated that other examples could have other configurations.

20

In this example, the PSU cartridge enclosure 145 is fabricated from steel sheet to form a housing portion that includes the bottom face 149, the side faces 150 and 151 and the front and rear faces 146 and 147. Cover portions that are secured to the housing portion form the top face 148 and top front face 148a. The cover portions are secured to the chassis portion by suitable fixings, for example one or more screws 152. It will be appreciated however, that in another example, other faces, or portions, of the enclosure could form the chassis and the cover portions. The provision of the enclosure 145 means that the PSU cartridge 81 can safely be handled by an operator who is not a skilled technician. Also, through the use of the enclosure 145, the PSU

cartridge 81 is a robust unit that protects its inner workings from the outside environment and vice versa. The use of a conductive enclosure, e.g., a metal enclosure, means that the PSU cartridge includes its own electromagnetic shielding. To this end the PSU enclosure 145 is provided with EMI fingers 153 to ensure good
5 contact with the shelf chassis and the adjacent components.

As shown in Figure 4, the PSU cartridge 81 incorporates two D-shaped handles 156 to facilitate insertion and removal of the PSU cartridge 81 with respect to an aperture 82 in the rear face of the shelf enclosure. A latch member 155 can be
10 pivotably mounted on a plate that can be secured (e.g., using screws) to the rear face of the shelf enclosure. The latch member 155 is configured to engage one of the handles 156 and to secure the PSU 81 in place. In other examples, the PSU 81 could be provided with an injector/ejector handle in a manner similar to the information processing cartridge. As shown in Figure 7, the front face 147 of the PSU cartridge 81
15 has perforations 157 to allow for airflow into the PSU cartridge 81. As shown in Figure 4, the rear face 146 of the PSU cartridge 81 also has perforations 159 to allow for air to be exhausted from the rear of the PSU cartridge 81.

A pair of fans can be located behind the perforated portions 159 of the rear
20 face of a PSU cartridge 81 to channel cooling air through the PSU cartridge from the front to the rear. LED indicators 161 can be provided on the rear face 146 of the PSU enclosure 81 to indicate whether input power is good, whether output power is good, whether service intervention is required and whether the PSU can be removed. Electrical connectors 163 can be provided at the front face of the PSU (i.e. on the face
25 that in use is inside the shelf enclosure 47) for connection to the shelf. The PSU 81 of the present example may suitably employ an SSI-MPS (Server Systems Interface - Midrange Power Supply) compliant right angle connector at the front face 147 of the PSU 81 to connect to the shelf 41. The power inlet 83 for each PSU 81 can incorporate a cable/connector retention mechanism (not shown) on the rear face 146

of the PSU to prevent accidental or malicious removal of the power input cord from the PSU 81.

In the present example, the shelf unit enclosure provides slots at the rear of the
5 shelf unit for two hot-swappable, AC input PSUs 81. To provide redundant power supply, both PSU slots are populated. More generally, $N+M$ power supplies can be provided, where N is the minimum number of power supply units needed to support the components mounted in the shelf, and M is a selectable number of 0, 1 or more power supplies to provide for redundancy. In the present example, $N = M = 1$.

10

It will be appreciated that one possible construction of the CSSP cartridge 71 has been described and that as for the information processing cartridge 43, other examples could employ other materials and/or constructions.

15 If the full complement of information processing cartridges or switches are not fitted to the enclosure, then blanking panels/modules (e.g., the blanking panels 44 shown in Figure 3) are fitted to all unpopulated positions to maintain the correct airflow and thermal characteristics of the shelf, a safe internal operating temperature for the shelf life expectancy of all shelf components, electromagnetic compliance
20 (EMC) containment and electrostatic discharge (ESD) containment.

It will be noted that each of the FRUs, such as the information processing cartridges 43, is advantageously contained in its own robust enclosure to facilitate EMC containment, ESD containment, handling, storage and transportation. Each
25 FRU can be configured as a 'sealed' unit in the sense that it can be configured not to have field or customer serviceable parts internally. The FRUs can be configured readily to plug into the shelf enclosure and to be hot swappable. The FRUs can be keyed to prevent incorrect positioning and insertion into the shelf enclosure and are arranged positively to be retained in the shelf by a latching/locking mechanism.

The examples of FRUs described above are not provided with removable media. In the present example, internal data storage is provided by 2.5" IDE 9.5mm or 12.7mm profile hard disk drive (HDD) devices mounted internally in each information processing cartridge 43 and in the CSSP cartridge 71. The drives are not considered as FRUs and are not hot-swappable disk drives in the present example, although they could be in other examples. In other examples, the information processing cartridges can be configured without internal hard disk drives.

10 The internal configuration of the shelf 41 and a midplane 171 contained therein is described in the following with reference to Figures 8A, 8B 8C and 8D, and Figures 9A, 9B and 9C.

Figure 8A is a schematic plan view showing the internal configuration of an example of a shelf 41 with the cover 61 removed. Figure 8B is a schematic perspective view from above the rear of the chassis portion 47 of the shelf enclosure with the field replaceable units removed. Figure 8C is a schematic perspective view from below the front of the chassis portion 47 of the shelf enclosure with the field replaceable units and the base 51 removed. Figure 8D is a schematic perspective view from the front and above a part of the base 51 of the shelf 41. Figures 9A, 9B and 9C are, respectively, front, top and rear views of the midplane 171. In this example, the midplane is, in use, mounted vertically within the shelf 41 extending across the width W of the shelf 41 at a position approximately half way between the front and the rear of the shelf 41.

25

The vertically mounted midplane 171 extends, in this example, across the shelf 41 and allows for the electrical interconnection of the FRUs. The various apertures in the front and rear faces 57 and 59 of the shelf 41, in combination with the midplane 171, can be provided with guides (e.g., rails 181) and keying e.g., offset connector

positioning for the insertion of the FRUs into the enclosure and midplane 171. The midplane 171 can be a double-sided, or multi-layer printed circuit board (PCB) assembly that can be mounted vertically in a rigid manner within the enclosure. It can carry connectors 175 on a front surface 172 for making electrical connection with
5 corresponding connectors 120 on the information processing cartridges 43. It can also carry connectors 177 and 179 on rear surface 173 for making electrical connection with corresponding connectors 141 and 163 on the CSSPs 71 and the PSUs 81, respectively. Conductive tracks (not shown) on and through the midplane 171 can be provided to interconnect the various connectors. In addition, the midplane can
10 provide connectors for receiving corresponding connectors connected to first and second indicator boards 183 and 184 that each carry a respective set of LED indicators 69. In the present example, the midplane 171 is not configured as a FRU and is not hot swappable. It is perforated to facilitate airflow through the shelf 41. The midplane 171 can include openings 185, which cooperate with openings in the
15 enclosures of the FRUs 43 and 81, to provide a path for cooling air to pass from the front to the rear of the shelf 41, the cooling air being driven by fans in one or more of the FRUs, for example in the PSUs 81, possibly also in the information processing cartridges 43.

20 A plenum chamber floor member 94 can extend horizontally from the front of the midplane 171 to the front face 57 of the shelf enclosure, or chassis 47. The member 94 provides a floor for a plenum chamber 66, which is supplied with air via the apertures 65 in the front bezel and, in the illustrated example, the slot shaped aperture 68 in the front face 57 of the shelf enclosure 47. Although, for reasons of
25 ease of illustration a slot shaped aperture 68 is shown, a plurality of apertures 68 aligned with the blade receiving locations may be provided. The aperture or apertures 68 can serve both as air vents for a flow of air to the plenum chamber 66, and also as latching locations for latching portions at the top of the injector/ejector levers 111 for

the blades shown in Figures 5B and 5C. The top and sides of the plenum chamber are provided by the top cover 61 and side faces 53 and 54 of the shelf enclosure 47.

- A plurality of cartridge guides 97 can be provided at the underside of the
- 5 plenum chamber floor member 94. In the present example, these guides comprise sprung wire members, e.g., of a resilient metal such as spring steel, that are attached to the top surface of the plenum chamber floor member 94 and extend through a plurality of apertures therethrough to result in a row of guides 97 at the underside of the plenum chamber floor member 94. This arrangement is shown in Figures 8B and 8C.
- 10 In Figure 8B, the sprung wire members 98 are shown attached to the top surface of the plenum chamber floor member 94. In the present example, the sprung wire members 98 are arranged in pairs, such that two guides 97 are provided by each spring clip 98. In Figure 8C, the guides 97 formed by the protruding portions of the sprung wire members 98 are shown at the underside of the plenum chamber floor member 94.
- 15 Each guide 97 is advantageously positioned so as to interface with the groove 108 in the plastics material in the upper face 104 of a processing cartridge 43 as shown in Figures 5A-5C to aid correct alignment and to facilitate insertion of the processing cartridge during insertion of the cartridge into the shelf 41. The use of the spring clip as a guide 97 also serves to urge the processing cartridge downwards to provide a
- 20 secure mounting of the processing cartridge 43, to take account of manufacturing and operational tolerances and to assist in insertion of the processing cartridge where an operator does not align this absolutely correctly.

- A further row of cartridge guides 99 can be provided at the upper surface of
- 25 the base 51 of the shelf 41. In the present example, as shown in Figure 8D, these guides 99 have a rail like form, which can be achieved by punching or stamping through the base 51 of the shelf 41. In this example each guide, or rail, 99 includes a pair of upstands separated by an aperture 100 through the base 51. The size of the aperture 100 can correspond to the width between the upstands. The separation of the

upstands is selected so that the overall width of the resulting rails is slightly less than the width of a groove formed in the lower face of an information processing cartridge 43. Thus, each guide 97 is advantageously arranged so as to interface with the groove 1113 in the plastics material in the lower face 104 of a processing cartridge 43 as shown in Figures 5A-5C to aid correct alignment and to facilitate insertion of the processing cartridge during insertion of the cartridge into the shelf 41.

In the present example, where the guides 97 and 99 are formed from metal, the provision of the grooves 108 and 113 in plastics material at the upper and lower faces, respectively, of each information processing cartridge 43 (see Figures 5A-5B) results in a combination of metal and plastics materials that can provide a low friction interaction, facilitating insertion of the information processing cartridges.

If, for example, the information processing cartridge enclosure is made of a metal, it may be undesirable to provide metal guides to avoid a metal to metal interaction. In such a case, for example, it may be desirable to form the guides from a plastics material having a low coefficient of friction, such as polytetrafluoroethene (PTFE) or polythene. Plastics rails could be attached to the underside of the plenum chamber floor member 94 and/or on the upper surface of the base 51 of the shelf 41. In such an example, grooves on the upper and lower faces of the information processing cartridges 43 could then be formed of metal or plastics and still result in a low friction arrangement.

A CSSP/PSU divider 96 can be provided to the rear of the midplane 171 and can extend horizontally to the rear face 59 of the shelf enclosure 47. The CSSPs 71, when inserted, are supported by the divider 96. To aid the correct insertion of the CSSPs 71, CSSP guide pins 178 are provided on the midplane 171 at positions adjacent connectors 177 on the midplane 171 for connection to the CSSPs 71.

Respective positions 88 and 89 can be formed in the front face 57 and the rear face 59 at which first and second indicator boards 183 and 184 supporting the indicator LEDs 69 can be located. These positions 88, 89 therefore include an aperture through the respective face of the shelf enclosure 47 such that indicator LEDs 5 69 mounted onto a circuit board attached to the inside of the shelf enclosure 47 may be viewed from outside the shelf enclosure.

There now follows are more detailed description of the midplane 171.

10 As mentioned above, the midplane 171 connects all the elements of a shelf together, including, in the present example, up to sixteen information processing cartridges 43, up to two CSSPs 71, two PSUs 81 and the two indicator boards 183 and 184. In the present example, due to its location within the shelf enclosure, the midplane 171 is not configured to be swappable. Accordingly, to maximize the 15 system reliability, the midplane is configured to provide as a high level of reliability as possible. To this end, the midplane is advantageously configured without active devices and to include the minimum number of decoupling capacitors consistent with good design practice (ideally zero).

20 The midplane supports a number of paths for various power and signal lines to interconnect the FRUs.

In the present example, each information processing cartridge 43 has a high speed information signal connection (e.g., a Gigabit (Gb) Ethernet 25 SERializer/DESerializer (SERDES) connection) to each of the CSSPs 71, each connection consisting of two pairs of differential signals. In a conventional manner therefore, the tracking of the paths for these signals is arranged to keep the pairs well balanced and on a single signal layer (i.e. without vias) to support such differential signals at high frequency.

In addition, in the present example, each information processing cartridge 43 has a serial console connection to the CSSP cartridge 71. Each connection consists of two TTL (Transistor-Transistor Logic) level signals that make a transmit and return
5 (TX and RX) pair.

Also, each PSU 81 has a management signal connection (e.g., a serial I2C (Inter-IC Bus) connection) to the CSSP cartridge 71 to control power and monitor environmental parameters. The I2C bus comprises of two signals SCL and SDL
10 (serial clock line and serial data line). In addition, an I2C address programming pin is provided for the PSUs 81.

Each information processing cartridge 43 and PSU 81 can signal to the CSSP cartridge 71 that it is inserted by pulling to ground (GND) a respective Inserted_L
15 signal (i.e., an active low signal). These signals are fed to the CSSP cartridge 71 via the midplane 171.

Each PSU 81 has five 12 Volt output rails. The routing from each PSU 81 is arranged so that a fault in any single FRU cannot completely interrupt the power to
20 any other.

As mentioned above, the midplane 171 is provided with appropriate connector arrangements for receiving the connectors on the FRUs.

25 In the present example, the information processing cartridge 43 connects to the midplane 171 through a 40 pin Single Connector Attachment (SCA-2) connector as defined by the Small Computer Systems Interface (SCSI) standard. Accordingly, the midplane carries corresponding connectors 175.

In the present example, each CSSP cartridge 71 connects to the midplane 171 through a two right-angle 20 pair connector (e.g., 2mm HM-Zd connectors available from Tyco Electronics). The corresponding connectors 177 on the midplane are straight male parts with a power connector. A guide pin arrangement is provided in addition to the connectors to prevent misaligned modules causing bent pins during insertion. The guide pin also provides a leading ground. The CSSP cartridge 71 also connects to the midplane 171 through a right-angled 125 way 5 row 2mm connector. The connector 177 on the midplane 171 includes a straight male part. A guide pin arrangement is provided in addition to the connectors to prevent misaligned modules causing bent pins during insertion.

In the present example, as mentioned above, each PSU 81 connects to the midplane 171 through an SSI-MPS specification connector. The contacts are configured 5P/24S/6P with sequenced signal (S) and power (P) pins. Where the connector on the PSU is a 1450230-1 R/A male header, solder tails connector, the mating connector 179 on the midplane can be a 1450540-2 vertical receptacle, press-fit connector.

In the present implementation, indicator boards 183 and 184 (see Figure 8A) are provided at the front and rear of the system and are configured as FRUs. In this example they hold three system-level indicator LEDs 69 and include a FRU identity (FRU-ID) programmable read-only memory (PROM) each. Three LEDs 69 are present on the indicator board. There can, for example, be a white locator LED that can be switched by the user for locating the system; a green power-on LED to indicate when the system is powered; and an amber service-required LED to indicate a fault or other condition requiring servicing. These LEDs can be driven by the CSSP 71.

In the present example, identification information (FRU ID) for the midplane 171 is held on an I2C electrically erasable programmable read only memory

(EEPROM) in the front indicator board 183. In addition to the I2C signals necessary to access the FRU ID EEPROM, the CSSPs 71 provide a current limited supply to the indicator boards 183 and 184 via the midplane. The indicator boards 183 and 184 are also provided with an I2C address programming pin. Depending on the
5 implementation, FRU ID information can be stored instead, or in addition, on the rear indicator board 184.

As the FRU-ID for the midplane 171 is held on one or both of the indicator boards 183 and 184, the midplane can be a totally passive unit. The FRU-ID PROMs
10 communicate with the CSSPs 71 via an I2C bus. Each device on the bus has a separate I2C address. The lower three I2C address bits of the EEPROMs used are available as pins on the device, to allow programming with resistors. The least significant bit of this address (A0) is passed to the midplane via the corresponding connector. This allows the midplane 171 to program the address of the FRU-ID
15 differently for the front and rear indicator boards 183 and 184, by pulling the address low for the front board and high for the rear indicator board 183. This ensures that both EEPROMS are available on the bus, at different addresses. The FRU-ID for the midplane can be stored on either front or rear EEPROM, but the present example the FRU-ID is stored in the EEPROM on the front indicator board 183. The EEPROM
20 can be 8kByte or larger.

As mentioned above, the midplane 171 includes openings 185 to provide a ventilation path for cooling air passing through the shelf 41. The cooling air passing through the shelf 41 via the midplane 171 can be driven by means of fans provided in
25 each of the information processing cartridges 43 and the power supply modules 81. The openings 185 shown in Figures 8B, 9A, 9B and 9C form schematic representations of openings in the midplane 171. In practice, the openings could have any form (i.e., a series of large openings, or a number of small perforations), arranged on the midplane to align with corresponding openings or ventilation apertures in the

various field replaceable units 43, 71 and 81. In this way, the path of the airflow from the front of the shelf to the back of the shelf can be configured to be as efficient as possible, depending on the detail configuration of the fan units and the ventilation openings or apertures in the information processing, switch, service processor and power supply unit modules 43, 71 and 81. Providing the fan units in the field replaceable units 43, 71 and 81, contributes to the aim of maintaining the chassis 49 and the midplane 171 of the shelf 41 free of active components, thereby minimizing cost, and facilitating maintenance. Also, by providing the fan units in each of the field replaceable units, merely inserting and removing field replaceable units automatically adapts the flow of cooling air to the number and type of field replaceable units inserted in the shelf 41.

As described above, in the present example each of the FRUs is designed to be a non-user serviceable unit. Thus each FRU presents the user with a "sealed" unit which may be inserted into and removed from the shelf 41 as desired or required. If a FRU ceases to be operable, then the user has a choice only of returning the FRU to a supplier or service company for repair or of discarding the non-operable unit. As the FRUs are non-user serviceable, there is no requirement for a skilled technician to be employed in inserting or removing the FRUs into or from a shelf 41. Thus each FRU is designed such that a non-skilled person should have difficulty in causing damage to the FRU during handling. Moreover, the configuration and construction of the FRUs (e.g., the injector/ejector levers, the grooves in the enclosures of the information processing units, etc), of the shelf enclosure and the midplane (e.g., the guide rails to guide insertion of the FRUs, the locating pins, etc) contribute to facilitating easy insertion and removal of the FRUs.

Shown in Figure 10 is an example of the flow of cooling air through the shelf 41 and FRUs 43, 71, 81 mounted therein.

In this example, the cooling air passing through the shelf 41 is drawn generally in a front to rear direction through the shelf 41 by cooling fans mounted within the CSSPs 71 and the PSUs 81. Two separate flow paths for cooling air are provided in this example. The first, indicated as flow path $\alpha\alpha$ by dotted lines 77 provides cooling
5 air to the CSSPs 71. The second path, indicated as flow path $\gamma\gamma$ by dotted lines 78 provides cooling air to the information processing cartridges 43 and PSUs 81.

The flow of cooling air along path $\alpha\alpha$ enters the shelf 41 through the aperture(s) 65 in the front face 57 of the shelf enclosure 47 into the plenum chamber
10 66. Further apertures (not shown) could also be provided in the side faces 53 and 55 of the shelf enclosure 47 to allow air to flow into the plenum chamber 66. This air then flows through the plenum chamber 66, and passes over the top edge of the midplane 171 to reach the perforations 133 of the front face of the CSSPs 71. The cooling air then passes through the CSSPs 71, providing cooling to the components
15 thereof before passing out of the CSSPs 71 through the perforations 135 in the rear face of the CSSPs 71 thus being exhausted from the shelf 41. This flow of cooling air along flow path $\alpha\alpha$ is driven by fans 79 mounted within the CSSPs 71. In the present example, a pair of fans 79 is provided within each CSSP 71 and is mounted against the rear face thereof.

20

Air flowing along path $\alpha\alpha$ is impeded from flowing around the processing cartridges 43 by plenum chamber floor member 94 and is impeded from flowing to the PSUs 81 by CSSP/PSU divider 96. This flow path $\alpha\alpha$ therefore ensures that air flowing to the CSSPs 71 is not warmed by passage through the processing cartridges
25 43 and therefore provides maximum efficiency cooling to the CSSPs 71.

The flow of cooling air along path $\gamma\gamma$ enters the shelf 41 through the perforations 115 in the front face of the information processing cartridges 43. The air thus enters the information processing cartridges 43 and provides cooling to the

components thereof. Cooling fans (not shown) within the information processing cartridges 43 direct the cooling air to the processor (CPU) of the information processing cartridge and direct the flow of air in the cartridge thereby increasing cooling efficiency. The air then exits the information processing cartridges 43
5 through the perforations 117 in the rear face thereof. The air then passes through the apertures 185 through the midplane 171 to reach the PSUs 81. This cooling air then passes through the perforations 157 on the front and upper front faces of the PSUs 81 to enter the PSUs and provide cooling to components thereof. It will be appreciated from Figure 10 that the sloping rear of the upper face of the PSUs 81 increases the
10 area over which air can be drawn into the PSUs, thereby reducing the back pressure on the air flowing through the shelf unit and aiding the cooling efficiency. The flow of cooling air along path $\gamma\gamma$ is driven by fans 85 mounted within the PSUs 81. In the present example, a pair of fans 85 is provided within each PSU 81 and are mounted against the rear face thereof.

15

Air reaching the PSUs 81 via path $\gamma\gamma$ will already have passed through the processing cartridges 43. Such air will therefore be already warmed above the ambient temperature outside of the shelf 41 by its passage through the processing cartridges 43. However, as the cooling requirement of the PSUs 81 is typically less
20 than that for the CSSPs 71, this does not cause any difficulty in the operation of the PSUs 81, which are adequately cooled by this flow of pre-warmed air. The pre-warmed air passing through the apertures 185 through the midplane 171 is impeded from flowing into path $\alpha\alpha$ and entering the CSSPs 71 by the SCCP/PSU divider 96.

25

As will be appreciated by the skilled addressee, the arrangement shown in Figure 10 is illustrative only and other arrangements whereby, for example, a mixture of cool air from the plenum chamber 66 and warmed air from the processing cartridges 43 is provided to each rear-mounted FRU can readily be constructed.

With reference to Figure 11, there now follows a description of functional elements of an information processing cartridge 43 as contained within the information processing cartridge enclosure 101.

5 The information processing cartridge 43 includes a microprocessor 192 (a non-limiting example of a microprocessor that can be utilized in the present example is an UltraSPARC™ processor). The microprocessor is mounted on an information processing cartridge motherboard 191.

10 A configurable clock generator 193, configured as a programmable clock synthesizer employing a crystal, can be used to produce CPU clock signals, CLKA and CLKB. The clock frequency can be determined by jumper settings (not shown). A vectored interrupt controller (I-Chip) 194 and a configurable core voltage regulator module (VRM) 195 are provided.

15 In the present example, memory means for use by the processor 192 when executing instructions can be provided in the form of buffered dynamic random access memory (DRAM), for example configured as dual in line memory modules (DIMMs) 196 with a 72-bit data path with error correction codes (ECC), seated in two sockets
20 on a riser card from the information processing cartridge motherboard 191. The memory capacity can be chosen to suit the processor addressable memory space. For example, in the present example, up to 4 Gigabytes (4GB) of addressable memory can be provided. Serial Presence Detect (SPD) auto-configuration is provided via a Service Management Bus (SMBus) over an I2C bus 197.

25 In the present example, a PCI bus architecture can be employed with a so-called SouthBridge bus bridge 199 with SuperIO and two Gb Ethernet Media Access Control (MAC) devices. As described above, however, other bus protocols (e.g., Infiniband) can be used. A 32bit PCI bus 198 can be provided from the

microprocessor 192. The SouthBridge 199 is a standard form of bus bridge, in the present example packaged in a 352 pin PBGA (Plastic Ball Grid Array) package, that provides the following functions: an SM Bus interface over the I2C bus 197 for access to the SPD (Serial Presence Detect) feature of the DIMMs that allows
5 initialization of the memory controller; an Xbus interface for access via an Xbus 200 (which is a packet switched multiprocessor bus) to a PROM 201, a real time clock (RTC) 202 and an information processing cartridge service controller (hereinafter termed a Blade Service Controller (BSC)) 203; an IDE (Integrated Drive Electronics) interface that provides an ATA-100 (AT Attachment) IDE connection 204 to an IDE
10 disk drive 205; and a serial console interface on a service bus 206 to the BSC 203 that is used for operating system functions including a console function with this embodiment.

For IO to the midplane 171, two AC-coupled Ethernet interfaces 207 and 208
15 are provided in the present example, which are packaged in a 316 pin PBGA. These Ethernet interfaces can provide a PCI attached Ethernet MAC capable of operation up to Gigabit Ethernet performance. The physical layer can be implemented using SERIALIZER/DESERIALIZERS (SERDESs) 209 and 210. An example of a SERDES device is the TLK2201 transceiver manufactured by Texas Instruments, Inc. The SERDES
20 devices use differential PECL TX+/- and RX+/- (Positive Emitter Coupled Logic Transmit and Receive) pairs to communicate to the switch portions of the CSSPs 71 over the midplane 171. The RX+/- pairs can be AC coupled at the information processing cartridge 43, the TX+/- pairs can be AC coupled at each CSSP 71. This facilitates hot-swap of the information processing cartridges 43 and the CSSPs 71.

25

Asynchronous serial connections 211 and 212 for communication between the BSC 203 and the Service Processor parts of the CSSPs 71 can be provided.

Internal data storage can be provided in the present example by a hard disk 205 with a capacity of 30GB or more rated for 24/7 continuous operation. The hard disk 205 is accessed using the primary IDE interface of the SouthBridge 199. The hard disk 205 can hold an operating system, for example a Solaris operating system, and
5 other software and data for performing information processing using the main, or host, processor (CPU) within the information processing cartridge 43.

In the present implementation, the BSC 203 can be implemented as a microcontroller (e.g., a Hitachi H8 microcontroller). The BSC 203 can provide
10 various functions, including for example: dual access (for the information processing cartridges and the CSSPs 71) to PROM 201 and EEPROM 213 for boot information and a FRU-ID for the information processing cartridge; channeling communication between an information processing cartridge 43 and the service processor part of the CSSPs 71; control of power on reset (POR), system reset and externally initiated reset
15 (XIR) to the microprocessor 192; control of the power, service-required and ready-to-remove LEDs 69; upgrading of field-upgradable firmware, via the serial interface; a watchdog function for the operating system; monitoring the speed of a CPU fan 214; and communications with an EEPROM 215 and the operating system via the Xbus 200.

20

In the present example, the BSC 203 can be powered by a 5V service bus (SB) rail as soon as a CSSP 71 and a PSU 81 are fully inserted into the midplane 171, it then turns on other DC/DC converters to provide power to the remainder of the information processing cartridge 43. A BSC reset signal can be derived from a simple
25 conventional power on reset (POR) generator that monitors a 5V supply rail.

In the present example a 1MByte Flash PROM 201 can be provided for storing boot variables for OpenBoot™ PROM (OBP) and Power-On-Self-Test (POST). Further OBP variables can be stored in a second 16kByte (16kB) I2C PROM 215,

- accessible via the SouthBridge SM Bus port over the IC Bus 197. The PROM 215 can contain 8kByte for OBP variables and 8kByte of unused space. A 16kByte I2C EEPROM 213 that is accessible via the BSC 203 can contain BSC variables and FRU-ID variables. The EEPROM is nominally divided into 8kByte for FRU-ID and 5 8kByte for the BSC variables. Write protection for the FRU-ID is implemented by BSC firmware. Such write protection may be carried out by, for example, acknowledging instructions to write to the protected area, but not to carry out those write instructions.
- 10 An environmental monitor sensor 215 can be provided to monitor the CPU and ambient temperatures. This sensor can be accessible via the onboard I2C bus from the BSC 203.
- The information processing cartridge 43 can be powered from two, diode 15 commoned, 9V power supply rails 216 and 217. DC/DC converters 218 can be used to provide the voltage levels required by the information processing cartridge 43. The DC/DC converters 218 are supplied by dual 9V inputs 216, 217, individually fused 219, 220 and then diode commoned 221, 222. A 5V DC/DC converter can be turned on as soon as the FRU is fully inserted, with the BSC 203 and required portions of the 20 SouthBridge 199 being powered (the 5VSB rail). A field effect transistor (FET) can be used to gate off the main 5V supply to the rest of the information processing cartridge 43. The DC/DC converter outputs and the main 5V FET can be arranged not to turn on until the BSC 203 turns them on via a signal from the SouthBridge 199. The SouthBridge 199 can be used so that if the BSC 203 is reset (by a watchdog 25 timeout or after a firmware download) the state of the DC/DC converters 218 is not affected. When the remaining outputs from the DC/DC converters 218 are within specification, a PWR_GOOD signal can be asserted low to the BSC 203.

A SouthBridge resume circuit can be operable to run from 3V3, and a simple Zener diode dropper circuit can be used to generate 3V3 from the 5VSB supply.

When the FRU is inserted the inrush current can be limited, for example to
5 <1A, and the rate of rise can be configured not to exceed a predetermined value (e.g., 20A/s) to provide a so-called soft start to facilitate hot-insertion. The intent is to prevent damage to the connectors and to avoid generating noise. A soft start controller 223, which controls a ramping-up of voltage levels, can be enabled when the predetermined signal (Inserted_L signal) is asserted low, this signal is on a short
10 pin in the connector and is connected to ground (GND – not shown) through the midplane 171.

In the present example, a processor impingement fan (processor fan) 214 is configured to run at full speed to cool the information processing cartridge 43 and the
15 fan. The speed of the processor fan and sink can be monitored by the BSC 203, using a tachometer sense pin on the microcontroller. In the event of the fan speed falling below a predetermined speed, or percentage of its nominal speed (e.g., 80%), the BSC 203 can be arranged to issue an alert. The nominal speed of the fan can be recorded as part of the BSC EEPROM contents.

20

The midplane connector 120 for the information processing cartridge 43 is used to establish the connection between the information processing cartridge 43 and the midplane. In the present example it supports up to 84 connections (pins) that will deliver SERDES outputs 224, 225, I2C signals 226, 227, and power 216, 217. Signal
25 connections may be made through a right-angled connector. Power connections may be made through the information processing cartridge right-angled connector. The connector can be configured to facilitate hotswapping of the information processing cartridge, for example by having a low insertion force and/or guide pins to increase the ease of serviceability and prevent module misalignment during insertion.

Interrupts to the processor 192 can be encoded using an encoded interrupt vector mechanism. An I-Chip Emulator (ICE) 228 functions as an interrupt concentrator, receiving all system interrupts and encoding them as an interrupt vector according to an interrupt vector code utilizable by the processor 192. In the present example, where an UltraSPARC™ processor is used, the interrupt vector encoding may be based on a 6-bit interrupt vector code.

With reference to Figure 12, there now follows a description of an example of a combined switch and service processor (CSSP) 71. In the present example, each CSSP 71 provides the functionality of a Switch 73 and of a Shelf Service Processor, or Shelf Service Processor (SSP) 74.

Figure 12 provides an overview of the functional components of the CSSP 71 including functional components of the Switch 73 and functional components of the SSP 74. In the present example, most of the components relating to the Switch 73 are mounted on a Switch PCB 231, and the components relating to the SSP 75 are provided on a SSP PCB 232. However, it should be noted that the components located in the lower portion of the switch PCB 321 (i.e., that portion below the SSP PCB 232 as illustrated in Figure 12 logically belong to the SSP 74, rather than to the switch 73. It will be appreciated that such component arrangements are not compulsory for successful operation and that any other component arrangement over any number of component boards can be easily achieved using conventional component arrangement techniques.

Firstly, with reference to Figure 12, there follows a description of functional elements of the Switch portions 73 of a CSSP 71 as contained within the CSSP enclosure 121.

- The midplane connector 141 on the CSSP 71 establishes the connection between the CSSP 71 and the midplane 171. In the present example, it supports up to 84 connections (pins) that will deliver SERDES outputs 265-268, I2C signals 310, 320, 321 and 322, and power 278, 279. Signal connections may be made through two 20-pair right-angled connectors. Power connections may be made through a right-angled connector. The connector can be configured to facilitate hotswapping of the board, for example with a low insertion force. The connector also uses guide pins to increase the ease of serviceability and prevent module misalignment during insertion.
- 10 A switch microprocessor 240 is provided, in the present example the microprocessor used is a Power PC (MPC8245) packaged in a 352pin Tape Ball Grid Array (TBGA) package. This microprocessor 240 supports between 1MB and 2GB of address space in the present example. It further includes an Embedded Programmable Interrupt Controller (EPIC) that provides 5 hardware interrupts (IRQs) or 16 serial
- 15 interrupts. There are 4 programmable timers with cascade mode function. DRAM memory for the processor can be provided in the present example by a commodity DIMM 242. The processor 240 can be connected to a 32bit PCI bus 241, which operates at, for example, 33MHz/66MHz.
- 20 A clock input to the processor 240 can be provided by a clock generator (CLK) 243. The CLK 243 can include a configurable clock generator (not shown) implemented as a programmable clock synthesizer employing a crystal used to produce CPU clock signals. The clock frequency can be determined by jumper settings (not shown). A vectored interrupt controller (I-Chip) (not shown) and a
- 25 configurable core voltage regulator module (VRM) (not shown) can be provided that operate substantially as described above with reference to the like components of Figure 11.

In the present embodiment two switch ASICs (application specific integrated circuits) 244, 245 are provided (in the present example, BCM5632 Gigabit switch ASICs). Each ASIC can provide twelve GMII Interfaces (1 Gigabit Ethernet) (for uplinks and downlinks) and one 10Gb XGMII interface for chip-to-chip
5 communication (bridging) 246 between the ASICs 244 and 245. Sixteen GMII 1Gb 'downlinks', in the form of serialized Gb Ethernet data, are provided through four quad SERDES 248-251 to allow each information processing cartridge 43 to communicate with the switch 73. Eight GMII 1Gb 'uplinks' are provided for external communication through two quad PHYs 253 and 254 (in the present example
10 BCM5404 ASICs) and RJ45 connectors on the rear panel 122. The ASICs 244 and 245 are configured via a PCI interface (32bit/33MHz) to the PCI bus 241.

A Flash PROM 256 can store a real time operating system, and management and configuration data for the microprocessor. The Flash PROM 256 in the present
15 example can be operable to hold 8MB - 16MB of data, depending on the software required. The flash PROM 256 can be operated via an on-chip XBus 258.

Also connected to communicate with the processor 240 via the XBus 258, a Real Time Clock (RTC) 259 can be provided for real-time functions with a back-up
20 battery.

Also connected to the XBus 258 can be a UART (Universal Asynchronous Receiver Transmitter) 260 which in turn connects to a serial bus 261 for providing an asynchronous console connection from the switch 73 to the SSP 74 which can be
25 accessed by the SSP.

An integrated MAC/PHY (Media Access Control/Physical) switch 271 can provides its own interface to the PCI bus 241. This MAC/PHY switch 271 can connects to a 10/100 Ethernet hub 272. The hub 272 can be operable to provide a

management interface to the SSP 74 and a connection from an external management network to the switch 73 and SSP 74 of a given CSSP 71. The connection from the integrated MAC/PHY device 271 to the SSP 74 can be coupled capacitively. A loopback mode can be provided by the MAC/PHY device 271 for system diagnostics.

- 5 The hub 272 can connect to an RJ45 connector 273 on the rear panel 122 of the CSSP enclosure 121.

- An 8kByte I2C EEPROM 262 can be used to store the FRU-ID and is accessible by the SSP portion 74 of each CSSP 71 via a serial bus 263 and the
10 midplane 171. The upper 2kByte of the EEPROM 262 can be configured to be write protected.

- An I2C Redundant Control Register (RCR) 275 can be used to provide an alternate, redundant path for powering-down the CSSP 71 and Shelf Level Indicators
15 69 mounted on the front 57 and rear 59 panels of the shelf 41. The I2C RCR 275 can be accessible by both the SSP 74 of the CSSP 71 containing the RCR and the SSP 74 of a further CSSP 71 connected via the midplane 171 via an I2C bus 276. In the present example, a device suitable for use as the RCR 275 is a Phillips PCF8574 IC.

- 20 With continued reference to Figure 12, there now follows a description of functional elements of the Shelf Service Processor (SSP) portion 74 of a CSSP 71 as contained within the CSSP enclosure 121 and provided on an SSP PCB 232.

- In the present example, communication between the Switch PCB 231 and the
25 SSP PCB 232 is facilitated by an interboard connector pair 298 and 299. It supports connections (pins) for I2C signals, 10/100 MAC/PHY output, and power. As described above, the switch PCB 231 carries the components associated with the switch, and it also carries the power, FRU-ID and environmental monitoring components along with the connectors for connections to the midplane 171 and

external connectors. Thus, in the present example, all SSP components requiring a connection to the midplane 171 or an external connection have signal paths routed through the connector pair 298, 299 and via the switch PCB 231 to the relevant midplane or external connectors.

5

In the present example, the SSP 74 includes a microprocessor 301 (e.g., a Power PC (MPC8245) processor) mounted on the SSP printed circuit board (PCB) 232. The processor 301 can be connected to a PCI bus 302, the present instance a 32 bit bus that operates, for example, at 33MHz/66MHz.

10

A clock input to the processor 301 can be provided by a clock generator (CLK) 303. The CLK 303 can comprise a configurable clock generator (not shown) implemented as a programmable clock synthesizer employing a crystal used to produce CPU clock signals. The clock frequency can be determined by jumper 15 settings (not shown). A vectored interrupt controller (I-Chip) (not shown) and a configurable core voltage regulator module (VRM) (not shown) can be provided that operate substantially as described above with reference to the like components of Figure 11.

20

The processor 301 can be provided with a DRAM memory 305. The memory capacity can be chosen to suit the processor addressable memory space. In the present example, 8 MB of DRAM memory is provided.

25

An integrated MAC/PHY switch 306 can provide its own interface to the PCI bus 302. The MAC/PHY switch 271 can be connected to 10/100 Ethernet hub 272 via the interboard connectors 298, 299. A loopback mode can be provided by the MAC/PHY switch 306 for system diagnostics.

Octal UARTs 308 and 309 can be connected between the PCI bus 302 and the interboard connector pair 298, 299. The signal path can be continued from the interboard connector pair 298, 299 to serial connections 310 on the midplane connector 141 on switch PCB 231. The Octal UARTS 308, 309 can facilitate serial
5 communications between the SSP 74 and each of the processing cartridges 43.

Also connected to the PCI Bus 302 can be a dual UART (DUART) 312 that in turn can connect via the interboard connectors 298, 299 to serial bus 261 for providing an asynchronous console connection from the SSP 74 to the switch 73. The DUART
10 312 can also have an I2C connection to an external connector on the rear face 122 of the CSSP enclosure 121. The external connector can provide a common operating system/boot console and command port 311.

Connected to the processor 301 via an XBus 314 can be a Flash PROM 315.
15 The Flash PROM 315 can store a real time operating system, and management and configuration data for the microprocessor 301. The Flash PROM 315 can be operable in the present example to hold up to 2MB of data, depending on the software required.

Also connected to the processor 301 via the XBus 214 can be a real time clock
20 (RTC) 316 for real-time functions with a backup battery. The RTC 316 can also provide 8kByte of non-volatile random access memory (NVRAM), in the present instance implemented as an EEPROM. This can be used to contain information such as the FRU-ID, a serial number and other FRU information.

25 To facilitate I2C communications between the SSP 74 and the other CSSP 71, the midplane 171 and the PSUs 81, a multiplexer 318 can be provided. The multiplexer 318 can have a single I2C connection to the processor 301 and connections, via the interboard connector pair 298, 299 and the midplane connector 141 to both PSUs 81, the midplane 171 and the other CSSP 71.

The processor 301 can also comprise an embedded DUART to provide a redundant serial link to the SSP 74 of the other CSSP 71. Although it would be possible to implement this link using an external DUART, the advantage of using an
5 embedded DUART is that the connection to the other CSSP is reliable and therefore likely to be functional. Where the embedded DUART link does not use the I2C Multiplexer for communications to the other CSSP, a common mode of failure for both the SSP – SSP I2C links can be avoided, it being assumed that the processor 301 is likely to be functional even if both embedded DUART channels are non-functional.

10

The CSSP 71 can powered from two, diode commoned, 9V power supply rails 278 and 279. DC/DC converters 281 can be used to provide the voltage levels required by the CSSP 71. The DC/DC converters 281 can be supplied by dual 9V inputs 278, 279, individually fused 285, 286 and then diode commoned 287, 288. A
15 soft start controller 283 can be provided to facilitate hot-insertion. A 5V DC/DC converter (I2C power regulator) 282 can be turned on as soon as the CSSP 71 is fully inserted. A 3.3V DC/DC converter can be turned on when instructed, for example through SSP service software, by asserting low an appropriate signal (ON_L – not shown). The 3.3V converter can be arranged to turn on a converted for 2.5V, 1.2V,
20 and a processor core voltage rail (Vcore) when the voltages are within an appropriate range.

When the CSSP 71 is inserted the inrush current can be limited, for example to <1A, and the rate of rise can be configured not to exceed a predetermined value (e.g.,
25 20A/s) to provide a so-called soft start to facilitate hot-insertion. The intent is to prevent damage to the connectors and to avoid generating noise. A soft start controller 283, which controls a ramping-up of voltage levels, can be enabled when the predetermined signal (Inserted_L signal) is asserted low, this signal is on a short pin in the connector and is connected to ground (GND – not shown) through the

midplane 171 until one of the supplies is removed. These circuits can be configured to withstand an overvoltage at their inputs whilst the input they are feeding is not powered, without any leakage to the unpowered circuit. A sense circuit can detect if the voltage has dropped below a threshold, for example 2.0V, as a result of a blown
5 fuse, a power rail going down, etc. The DC/DC converters 281 can be protected against short circuit of their outputs so that no damage occurs.

The I2C regulator 282 can be powered as soon as the CSSP 71 is fully inserted into the midplane 171. This can be facilitated through short pins connected to the soft
10 start controller 283, which controls a ramping-up of voltage levels. The other DC/DC regulators can be turned on, for example by SSP software.

A pair of fans 290, 291 can provide cooling to the CSSP 71. The fans 290, 291 can be configured to run at full speed to prevent overtemperature conditions by
15 minimizing the temperature of the internal components and the fan. The speed of the fans 290, 291 can be monitored by the SSP 74 through an environmental monitor 295 on the switch board 231. The environmental monitor 295 can be alerted in the event of the fan speed falling below a predetermined value (e.g., 80% of its nominal speed). The fan can provide tachometer outputs to facilitate the measurement of fan speed.

20 LED indicators 137 can be provided, for example with a green power LED, an amber LED for indicating that service is required and a blue LED for indicating that the switch is ready to be removed. LED indicators integrated on 2x4 stacked RJ45 connectors on the rear face of the CSSP 71 can be arranged, for example, to show
25 green continually when the link is present and flash green when the link is active.

The environmental monitor ENV MON 295 can be provided to maintain operational integrity of the CSSP 71. The ENV MON 295 can include limit values in limit registers and can monitor, for example, temperature within the CSSP enclosure

121, the CSSP power rails, including the 12V, 3V3, Switch Processor Core Voltage, CSSP Processor Core Voltage and the two 9V power feed rails 278, 279 from the midplane 171. The outputs of the DC/DC converters 281 can be fed in to A/D inputs of the ENV MON 295 for Watchdog comparisons to be made to the voltage limits set
5 in the limit registers. As noted above, the ENV MON 295 can also monitor the operating speeds of the fans 290 and 291. The ENV MON 295 can communicate with the SSP 74 of both CSSPs via an I2C bus 296.

For IO to the midplane 171 shown in Figures 8A-10, the midplane connector
10 141 can include sixteen 1Gb Ethernet connections 265-268 from four quad SERDES 248-251 and the I2C bus lines 596.

The SSP 74 can access the I2C devices (FRU-ID EEPROM, 8-bit I/O expansion chip, and the system hardware monitor) through the midplane 171.
15

For external IO, rear panel Gb Ethernet connections can be provided from the two quad PHYs 253, 254 to 2x4 stacked RJ45 connectors 139 (to give 8 uplinks). Each port can be an independent 10/100/1000 BASE-T (auto negotiating) port. The PHY devices 253, 254 can operate in GMII mode to receive signals from the 8-
20 Gigabit interfaces on the ASICs 244, 245.

The Power Supply Units (PSUs) 81 can configured such that when two or more PSUs 81 are connected in parallel in the shelf 41, failure of any one of the paralleled units shall not affect system operation. Moreover, one of the PSUs can be
25 installed or removed from a “live” system with or without input power applied. The outputs can have overcurrent protection.

The PSU can have an I2C interface to provide power supply status via the midplane 171. The PSU can have an internal temperature sensor that reports via the

I2C interface. The PSU fan speed can also be monitored and errors are reported via the I2C interface. Overvoltage and overcurrent sensors can also report via the I2C interface.

5 There now follows a description of aspects of an example of a power supply 81 with particular reference to Figure 13.

When a power supply (e.g. mains power, or UPS type protected power) is connected to the cable connector 83, transformer, regulator and rectifier circuitry 400
10 can operate to generate a DC output (in the present example, 12V DC) from the input (in the present example 230/240V 50Hz AC or 110V 60Hz AC).

In order to ensure good cooling reliability within the shelf 41, each PSU 81 can have a pair of cooling fans 402, 403 located at the rear of the PSU enclosure as
15 described above with reference to Figure 7. As a failsafe measure against a failure of one PSU 81, the fans of each PSU 81 can be powered by both PSUs 81. Thus, in the event that one PSU 81 is non-operative to produce the DC supply for the shelf 41, as long as the other PSU 81 is operative, not only will all components of the shelf 41 be powered as normal, the fans of both PSUs 81 can continue to run.

20

As shown in Figure 13, this dual powering of cooling fans 402, 403 can be effected by providing a power supply line 404 from the transformer, regulator and rectifier circuitry 400 to power both fans 402, 403. Also, first and second separate power lines 410, 412 from the other PSU 81 can provide duplicate power supply to
25 the first and second fans 402, 403 respectively. The fan 402 can thus be powered by a diode commoned supply from line 404 and a diode commoned supply from line 410. Diode protection can be provided by diodes 405 and 411 respectively. The speed of the fan 402 can be controlled by a speed controller 408. Similarly the fan 403 can be powered by a diode commoned supply from line 404 and a diode commoned supply

from line 412. Diode protection can be provided by diodes 406 and 414 respectively. The speed of the fan 403 can be controlled by a speed controller 409.

The two speed controllers 408, 409 can in turn be controlled by a data input
5 from each CSSP 71 received via an I2C bus connection (not shown in Figure 13). Power supply lines carrying DC power for the other FRUs of the shelf 41 are shown in Figure 13 as power line 416. All power connections to and from the PSU 81 can connect to the midplane 171 when the PSU is inserted in the shelf 41 via the midplane connector 163. In the present example the PSU 81 connects to the shelf through a
10 5P/24S/6P configuration SSI-MPS compliant right angle connector 163 at the front face 147 of the PSU 81. Connectors for the I2C interface can also be provided.

The input power line 410 and 412 for each fan 402 and 403 can be provided with a softstart module 4131 and 4132 respectively, to allow for hot insertion of the
15 PSU 81 into the shelf 41. The softstart modules 4131 and 4132 can be controlled, for example, by pulling a signal to ground (e.g., a "mated" input line 4151 and 4152).

Where the two input power lines 410 and 412 are separate lines having separate softstart provision, there is no common failure mode for the backup method
20 of powering the fans 402, 403. Thus even if the PSU 81 ceases to be operable to generate the DC supply, and a component (power line or softstart module for example) fails in the supply path from the second PSU 81 to the fans 402, 403, at least one of the fans 402, 403 can remain operational as the shelf 41 still receives the cooling effect of three PSU fans.

25

In the present example, the power supply has four rear panel LED indicators 137. A blue "Ready to Remove" LED can be driven by the I2C interface and indicate that the power supply may be removed from the system. An amber "Service Required" LED can be driven by the I2C interface and indicate that the power supply

is in a fault condition: any output out of range, over-temperature or shutdown. A green "DC Output-OK" indicator can be driven by internal power supply circuits and show that the main 12 volt supply is functioning. The LEDs can remain lighted when individual outputs are in the current limited mode of operation. A green "AC Input-OK" indicator can be driven by internal power supply circuits and show that AC input power is within normal operating range.

With reference to Figure 14, there will now be described an example of data connectivity between the FRUs and midplane of the shelf 41. Power transmission paths are not illustrated in Figure 14. However, it will be appreciated that to facilitate the maximum component redundancy of the shelf 41, each PSU 81 can independently provide power to each FRU.

In the present example each of the processing cartridges (blades) 43 connects to the midplane 171 via a pair of information signal connections (e.g. Gb Ethernet links) 224, 225 and a pair of serial management signal connections 226, 227. Connections within the midplane 171 can ensure that each Ethernet link 224 is directed to a connection 265-268 from the midplane 171 to a first switch 73, and that each Ethernet link 225 is directed to a connection 265-268 from the midplane 171 to a second switch 73. Thus one Ethernet link can be established between each processing cartridge 43 and the switch 73 of each CSSP 71. Further connections within the midplane 171 can ensure that each serial connection 226 is directed to a connection 310 from the midplane 171 to the first SSP 74 and that each serial connection 227 is directed to the second SSP 74. Thus one serial link can be established between each processing cartridge 43 and the SSP 74 of each CSSP 71. As mentioned earlier, information signal connections other than Gb Ethernet connections (e.g. Infiniband connections) could be employed in other examples.

A plurality of serial connections can connect each SSP 74 to the other. Serial lines 320, 321 can connect each SSP 74 to the midplane 171 and connections within the midplane can connect the two sets of lines together. To provide a control interface from the SSPs 74 to the PSUs 81, serial lines 322 can connect each SSP 74 to the
5 midplane 171 and connections within the midplane 171 can connect to serial lines 324 from the midplane 171 to each PSU 81.

An example of data and control connectivity of the shelf 41 to and from computer systems external to the shelf 41 when the shelf 41 is arranged for use within
10 a multiprocessor server system such as that described above with reference to Figure 1 will be described with reference to Figure 15.

As summarized above with reference to Figure 14, in the present example each processing cartridge, or blade, 43 is connected to the switch 73 of each CSSP 71 by an
15 information signal connection (e.g. a 1Gb Ethernet link) formed by a combination of links 224, 225 from the processing cartridge 43 to the midplane 171, connections within the midplane 171 and links 265-268 from the midplane 171 to each switch 73.

Further, in this example a set of serial management signal connections
20 comprising links 320, 321 and connections within the midplane 171 connect the SSP 74 of each CSSP 71 to the SSP 74 of the other CSSP 71.

To provide external data connectivity between the shelf 41 and an external core data network 330, in association with which all information processing
25 performed by the processing cartridges 43 of the shelf 41 is undertaken, connections 331 can be formed between the core data network 330 and the eight 1Gb Ethernet ports 139 provided on the rear panel 122 of the CSSP enclosure 121.

In the present example, the connections by means of which control and configuration of the shelf 41 are performed are entirely separate to the connections to the core data network 330. Therefore, a first external switch 335 can connect to a management (I2C) port 273 of the first CSSP 71 and a second external switch 336 can
5 connect to a management (I2C) port 273 of the second CSSP 72. As described above with reference to Figure 12, the management port 273 can provide a management network interface to both the switch 73 and SSP 74 of each CSSP 71. The external switches 335, 336 can each be connected to each of a pair of System Management Server (SMSs) 338, 339. The SMS is not essential to the operation of the shelf 41,
10 but use thereof aids optimal operation of the shelf 41. In a typical multiprocessor server system a plurality of shelves 41 may be connected together via the core data network 330 under the control of a single management network utilizing one set of SMSs 338, 339. A set of SMSs 338, 339 may comprise a single SMS (as well as a plurality thereof). However use of at least two SMSs enables redundancy of
15 components, therefore increasing overall system reliability.

A serial interface control 343 operable under telnet protocol control is also connected to the shelf 41 in the present example. This can provide a common operating system/boot console connection to the SSP 74 of both CSSPs 71 via the
20 RJ45 connector 311 on the rear panel 122 of each CSSP enclosure 121.

It will be appreciated from the above that a flexible and scalable modular computer architecture has been described. In the described example up to 16 information processing cartridges, or blades 43, can be configured as sealed FRUs on
25 a single shelf 41, the number of blades being chosen according to customer requirements. Each blade has its own processor and random access memory. If, for example, there is a maximum of 2Gbytes of memory per information processing cartridge, and one processor per blade, 16 processors (16P) with 5.33 processors per unit height (1U) and a total of 32GB of memory per shelf can be provided.

In the present example, the shelf 41 incorporates redundant combined switch and shelf service processor modules (CSSPs) 71 and redundant power supply units (PSUs) 81 separate from the blades 43. As the power supplies are carried by the shelf, the information processing cartridges can be kept compact and inexpensive. Also, as a result, they can be powered by DC power only, via the midplane 171.

Also, as mentioned earlier, the FRUs (e.g., the information processing cartridges, or blades, 43, the CSSPs 71 and the PSUs 81) can all be configured as sealed units that do not contain any internal FRUs themselves and do not contain user serviceable items. The enclosures of the FRUs can be arranged to enclose all of the functional components of the FRU with only electrical connectors being externally accessible and with indicator LEDs being externally visible as well.

These factors can all contribute to keeping the cost of the FRUs low, as well as that of the overall system. The modular approach with the use of sealed modular field replaceable units for providing system functionality and with non-field replaceable units designed with a minimum possible number of active components enhances reliability. Moreover, easy and rapid maintenance is facilitated in the event of a failure of a FRU by simple replacement of that FRU, further reducing the cost of ownership.

Thus, it will be appreciated from the above description that the provision of a rack mountable shelf, that includes power supplies, a shelf service processor and switches in modular units, for carrying a number of processing cartridges, wherein the number of processing cartridges can be chosen according to customer requirements, provides a flexible and scalable computer configuration. The balancing of the load between the processors of the processing cartridges can be effected by software using conventional principles.

A configuration as described provides an easily scalable processor architecture, whereby the processing power provided by a complete system based on the information processing cartridge/information processing cartridge carrier
5 architecture can be scalable from moderate to very high capacity through the simple addition of further information processing cartridges.

As illustrated in Figure 16, an example of the external connections from a shelf 41 can be in the form of two active information signal connections (e.g.,
10 Ethernet connections) 350 and 351, two active power connections 353 and an active/standby pair of management connections 354. With regard to the management connections, each connection comprises a serial connection and a network (e.g. Ethernet or Infiniband) connection. It is possible to connect to either the active or the standby connection, as the incoming signal will be internally routed to whichever
15 management controller (CSSP) is the current master. It will be appreciated, therefore, that the connections to a shelf can be kept to a minimum. It will further be appreciated from the configuration shown in Figure 16 that the system is scalable beyond a single shelf unit 41.

20 Figure 17 illustrates how a plurality of shelves can be configured within one (or more) racks to provide even higher processing power. Such a constellation of shelves to provide a large grouping of servers is sometimes termed a “web farm” or “server farm” 360. As shown in Figure 17, the web farm comprises a plurality of shelves 41 that each carry a plurality of blades 43. Also provided are a plurality of
25 Network Attached Storage devices (NAS) 373 for providing storage for critical data, e.g. email data storage, for the web farm. The NASs 373 are not required if there is no critical data to be stored, e.g. if the web farm is operating solely to provide web caching services.

Management control of the web farm 360 can be provided through a pair of System Management Servers (SMSs) 362. Each SMS 362 can be connected to a management network via a link 366 and to a management console 365. The SMSs 362 can communicate with the individual shelves 41 via a pair of management switches 364. Each shelf 41 and NAS 373 can be connected to each management switch 364 via a connection 367. Thus dual redundant management connections can be provided to each shelf 41 and NAS 373.

Flow of data to and from the web farm 360 can be provided through a pair of data switches 369. Each data switch 369 can be connected to a consumer network via a link 370. It is to be understood that the consumer network can be a larger data network to which the web farm 360 is connected. This network can be an office or corporation intranet, a local area network (LAN), a wide area network (WAN), the Internet or any other network. Connections between the data switches and the shelves 41 can be facilitated by connections 371. It is to be noted that as each shelf has its own switching capability, there is no need for each shelf 41 to be directly connected to the data switches 369. Connections can also be provided to connect the NAS units 373 to the shelves 41. The topology used for interconnection of the data switches 369, shelves 41 and NASs 373 can be any topology providing at least one connection of any length between every possible pair of units. Complex topologies arranged to minimize the maximum connection length between any two given units in the web farm can be used.

The web farm 360 comprising a plurality of shelves 41 with or without a plurality of NASs 373 can suitably be used as any or all of the entry edge server group 9, web edge server group 15 and application servers 19 described above with reference to Figure 1.

As an alternative to providing critical data storage within a NAS 373, such storage can be provided within one or more NAS cartridges fitted into one or more of the shelves 41 in place of processing cartridges 43. Another alternative is to provide a server shelf with local storage (such as a RAID array (Redundant Array of
5 Inexpensive Disks) in place of the NAS 373.

Thus there has now been described an example of a fully configurable computing system based on a plurality of self contained field replaceable units (FRUs) and scalable from a single processing cartridge with power supply and switching
10 capability to a multiply redundant multiprocessor server system with full system management capability extending over a number of co-operably connected server shelves. It will of course be readily apparent to the skilled reader that many of the specific features specified in the above description are in no way limiting and a variety of alternatives may be produced using only ordinary skill and common general
15 knowledge. Non-limiting examples of example modifications which may be made to the above described system are discussed hereafter.

There is no limit placed on the processing cartridges as to what software they should run. Each module within a shelf or farm may run under the same operating
20 system, or a plurality of different operating systems may be used. Examples of possible operating systems include Sun Microsystems' Solaris® OS or another UNIX™-Type OS such as Linux™, MINIX™, or Irix™, or UNIX™ or a Microsoft OS such as Windows NT™, Windows 2000™, Windows ME/98/95™, Windows XP™.
25

It is also not necessary that each processing cartridge within a shelf or farm be configured to run the same program software. For example, individual processing cartridges may be configured to execute, for example, fileserver software, mailserver

software, webhosting software, database software, firewall software, or verification software.

Although it has been described above with reference to Figures 4, 8, 12, 13
5 and 14, that functionality of a switch and of a shelf service processor should be provided within a single combined switch and service processor unit, this is not essential and separate switch and shelf service processor field replaceable units may be used.

10 Although it has been described above that a pair of PSUs and a pair of CSSPs may be provided so as to enable dual-redundancy, further PSUs and CSSPs may be provided so as to increase FRU redundancy further, thus providing statistically higher reliability.

15 In the power supply circuitry in each of the blades and CSSPs, two voltage sense circuits may be provided after the fuses and before the diodes, to prevent a latent fault caused by a failed fuse going undetected until one of the PSUs is removed or taken offline. Such circuits may be configured to withstand an overvoltage at their inputs whilst the input they are feeding is not powered, without any leakage to the
20 unpowered circuit.

Although it has been described above with particular reference to Figure 11 that the processing module may be based on an UltraSPARC™ processor, this is not limiting and any other processor having sufficient processing capacity to undertake the
25 tasks required of a particular processing cartridge may be used. Alternative processors include, but are not limited to, Intel x86 series and compatible processors, AMD x86 compatible processors, Alpha processors and PowerPC processors. The particular example of an x86 compatible processor is described in more detail with reference to Figure 18. In Figure 18, the parts corresponding to those of the UltraSPARC™ based

system of Figure 11 have the same reference numerals and will not be described again here. In a system based on an x86 compatible processor, the processor 378 itself communicates with the other components, including the memory 196 and PCI bus 198 via a Northbridge 379. The Northbridge 379 also includes an interrupt controller, so
5 no separate interrupt concentrator is required. The other components of the processing cartridge could be substantially the same as for the UltraSPARC™ based system described above.

Although it has been described above that each information processing
10 cartridge comprises a single microprocessor, this is not a limiting case as each or any of the information processing cartridges may have more than one microprocessor arranged to share common storage resources to operate synchronously (in lockstep) or asynchronously. Also, it is not necessary that all information processing cartridges inserted into a shelf at a given time are identical, rather a variety of different blade
15 architectures may be used simultaneously.

The provision of the functions of both Switch and Shelf Service Processor within a single FRU in the present example provides a facility within a single shelf 41 for dual redundancy in both functions in fewer different FRUs. As will be
20 appreciated, there is no restriction that these two functions must be provided within a single FRU and division of the two functions into separate FRUs would present no difficulty to the skilled addressee.

The backup power provision for the cooling fans of the PSUs is described
25 above with reference to Figure 13. Although it is described that the backup power supply to each fan should be independent of the other, if the risk of common mode failure of backup power supply is judged to be low, or is of low importance, then the backup supply to each fan may be provided in common with the backup supply to all other fans.

Also, an information processing module forming a field replaceable server blade can include a processor and memory can be configured by means of software, firmware or hardware to provide a special purpose function. By way of examples
5 only, an information processing module can be configured to perform the function of one or more of a firewall, or a load balancer, encryption and/or decryption processing, an interface to a secure network, e.g. a virtual private network (VPN), a specialized switch with wide area network (WAN) connectability.

10 Also, a storage blade may be provided. The storage blade can be configured to be mountable in a server blade receiving location in a blade server carrier. The storage blade can comprise storage blade connectors configured for interconnecting with carrier connectors on the server blade carrier, whereby the storage blade is interchangeable with a server blade. A carrier, or shelf, for a server system can be
15 arranged with a plurality of blade receiving locations for receiving blades, wherein the blades can be storage blades or information processing blades. The server system can be self configuring on receipt of the blades according to the type of blade received in each said location. To achieve this the blade service controller in each blade can be operable to communicate with a shelf service processor to perform said configuring.

20 In operation of a multiprocessor system such as a blade server system incorporating a plurality of server blades as described above, a plurality of domains will be established for each of the blades 43 and for the CSSPs 71. In such an environment where multiple domains are provided, there is a need for a central
25 repository of allocated unique host identities for the servers. Depending on the particular implementation, the host identities could be in the form of so-called "host IDs", simply a media access control (MAC) address for each of the servers, or some other identity that relates to the individual server domains. One of the reasons for

providing unique host identities for the individual servers relates to the association of software licenses with the individual servers.

If an individual CSSP or server develops a fault, and has to be replaced, then
5 the replacement server should receive the same host identity in order that it can resume providing the same services without the need for network configuration. By assuming the same identity, apart from other reasons, it is then possible for the software licenses to be transferred to the replacement processor.

10 In the following, there will be described an example of the present invention in which host identities can be stored in a secure manner within a multi-server, multi-domain system, such as the blade server system described above. A host identity (hostID) is a unique identifier that can be assigned to only one processor domain of a server or other system module at a time to uniquely identify that processor domain.
15 Accordingly, the host identity could also be termed a host identifier. Where there is only one processor domain in a server, then the host identity can also be used to identify that server. A host identity can then be used, for example, for registering licenses for software operating on that processor domain, or a server, etc. As indicated in the introduction, the term "host system" is used herein to define a processor
20 domain, or a server, processor, or module having one processor domain, such that a host identity can only be active for one host system at a time. Also, although the term "host identity" is used herein, it will be appreciated that other terms could be used instead for such an identity, such as, for example, only, a network identity, a network address, a system address, etc.

25 As indicated in the introduction, where one or more of a set of available system resources can be allocated to perform a processing task, it may be desirable to reallocate the system resources to carry out the processing tasks.

For example, if the system resource(s) originally allocated to perform the processing task fail(s), then it may be desirable to allocate one or more different system resources to enable further processing for that processing task. In view of the software licensing aspects identified above, it is important that an appropriate
5 reallocation of the host identities takes place when system resources and/or processing tasks are reallocated in such a failover situation. For example, if a blade server in a system as described above were to fail, it may be desirable to allocate a "spare" blade server from a pool of one or more such spare blade servers. In such a case, the host identity allocated to the failed blade server would need to be transferred to the "spare"
10 blade server while still ensuring that a host identity is allocated to only one active host at a time.

Also, the transfer of host identities may be desirable to enable dynamic upgrading or modifying of the multiprocessor system. For instance, during use, an
15 initially expected hardware resource requirement of a particular user may prove to be insufficient. It will be appreciated that, particularly in an environment where multiple users are sharing the overall resources of the system, it would be undesirable to have to reinitialize the whole system in order to accommodate changing requirements of one or a selection of the users.

20

An embodiment of the present invention as described in the following thus provides a mechanism enabling host identities to be reallocated between host systems, for example between blade servers or service processors, in a secure manner while still ensuring that a host identity is allocated to only one host system at a time.

25

In the present example, as described with reference to Figure 19, the approach is taken to store allocated host identities for the service processors 71/74 and the server blades 43 in the shelf 41 and for copies of those host identities to be stored in two of the field removable components within the server blade system. In the

particular example to be described, one copy of the host identities can be stored in each of the two CSSPs 71. Although, in the present instance, two copies of the host identities are stored in field replaceable components of the blade server system, it will be appreciated that another number of sets of copies could be stored instead.

5

Figure 19 illustrates an example of the present invention where the host identities allocated to a server blade shelf 41 can be stored in non-volatile memory 411 (for example an EEPROM) in the front indicator board 183. This can be the same non-volatile memory 411 that is used in the front indicator board for storing the FRU ID for the midplane 171. A set of host identities 410 for components within the shelf that, such as the CSSPs 71 and the server blades 43 can be stored in the non-volatile memory 441 in the present instance. In a particular example, the host identities can be host IDs which are used for controlling software licensing and identifying the processor domains to the network to which the blade server system is connected.

15

Optionally, a copy 412 of the host identities 410 could be stored in non-volatile memory 413 (for example an EEPROM) in the rear indicator board 184.

In the present example, the host identities that are stored include a host identity 20 for the first CSSP (CNI0), a host identity for the second CSSP 71 (CNI1), and a host identity for each of the up to sixteen blades 43-0 - 43-15 (BNI0 - BNI15) that may be received within the shelf.

In normal operation of the blade server system, a copy of the host identities 25 410 (CNI0 - BNI15) is stored in each of the CSSPs 71. In the example illustrated, a first copy 420 is stored in a registry, or memory, 421 within the first CSSP 71-0. A copy 430 of the host identities may also be stored in a registry, or memory, 431 in the second CSSP 71-1. Also, in normal operation, one of the CSSPs 71 acts as a master CSSP, with the other of the CSSPs serving as a back up. In the present example, if

both of the CSSPs are present and functioning correctly, it will be the CSSP that occupies slot 0 that assumes the role of master. However, it will be appreciated that this is merely one example, and the CSSP occupying slot 1 could instead act as the master, or indeed this could be determined in a random manner.

5

The CSSP 71-0 that acts as the master is then responsible for distributing the host identities to the individual domains within the blade server system, specifically to each of the blades which are occupying blade receiving locations within the blade server system, and to the second CSSP1 if present. Within each of the blades 43-0 to 43-15 that are present in the blade system, the host identity for that blade 440-0 to 440-15 in respective registries, or memories, 441-0 - 441-15, respectively. The memory in which the host identities are stored in the CSSPs 71 and in the server blades 43 will typically be non-volatile memory such as an EEPROM.

15 As mentioned above, Figure 19 illustrates the normal operation of the system. Although a copy of the host identities is stored in each of the CSSPs, only one host identity per host system is active at any time. In the present instance, the CSSPs 71 and the respective blades 43 each form respective host systems. The active host identities are the host identity CN10 for the first CSSP, the host identity CN11 for the second CSSP 71 and the host identities BN0 - BN15 (440-0 to 440-15) for each of the 16 blades 43-0 - 43-15 as held in the respective memories 441-0 - 441-15.

25 However, it may be in operation that it is desired to reorganize the distribution of the host identities. Typically, this could be as a result of identifying a failure of a host, for example if the SSP 74 of a CSSP 71 were to detect the failure of a server blade 43 due to a hardware or software fault. Alternatively, it could result from a desire to reorganize the host identities within a shelf, or between shelves, for example to take advantage of different models of blade having different performance levels.

In an embodiment of the present invention, one or more spare or additional host identities are held, which can then be provided for allocation to a host that needs to be configured to perform some processing tasks. The use of spare host identities means that downtime can be minimized when configuring a new host, although it does
5 have the disadvantage that additional software licenses are needed to be associated with the spare host identities.

Consider a situation where processing of a task is being performed on a host and then a fault develops with that fault. It would be desirable to continuing
10 processing of the task while the fault is corrected, for example by replacing the faulty host. If a spare host identity is available that has associated with it the appropriate software licenses to continue processing the task, then a different host could be configured using the host identity and the processing of the task could then be transferred to the new host. The old host could then be taken out of service and the
15 host identity associated therewith could then be used as a spare host for use in the event of future failover situations.

It would be possible for each master CSSP 71 (or SSP 74) to hold one or more spare host identities that could be used for configuring additional or new hosts in such
20 a situation. Consider then a situation where, for example, a server blade 43 that is responsible for processing a given task or tasks develops a fault. If a blade receiving location in the shelf were free, then a spare server blade 43 could be configured using a spare host identity from the CSSP 71, which spare host identity has associated with it the appropriate software licenses. The spare server blade could then continue
25 processing the given task or tasks. The CSSP 71 could then be operable to ensure that the faulty server blade was taken out of service and the host identity that was used by that faulty server blade is made available as a spare for future use. Typically, where the server blade has failed, the host identity will already then be inactive with respect to that server blade and that host identity can therefore be available for future use as a

spare. Alternatively, if the server blade is still operational, the CSSP 71 can be operable to issue a shutdown command to the server blade to cause it to shut down and thereby to relinquish the host identity. The master CSSP 71 can then hold the host identity as an available spare. It will be appreciated that this approach can

5 minimize the downtime with respect to processing of the task or tasks concerned, but that it does require the provision of additional host identities. The additional host identities could be held at the level of the CSSP 71. However, it may be desirable to hold the spare host identities as the level of a system administrator system to reduce the number of additional host identities (and associated additional software licenses)

10 needed, and also to facilitate the management of the allocation of host identities and software licenses.

Examples of operation in accordance with the present invention are described in the following, in which it is desired to move a host identity from one host system to

15 another.

Figure 20 provides an overview of components of the blade system described above useful in describing an example of the present invention.

20 Figure 20 illustrates an administrator system 450 that is connected via links 456 to a first master CSSP 71-A and a second master CSSP 71-B, the first and second CSSPs being in different shelves within the blade system.

The administrator system 450 could be configured as, for example, the

25 management console 365 illustrated in Figure 17 or could be a separate system connected via, for example, the management console 365 to the components of the system illustrated in Figure 17. The links 456 represent the connection paths between the administrator system 450 and the CSSPs 71-A and 71-B in the respective shelves. For example, if the administrator system 450 corresponds to the management console

365, the links 456 represent the connections through the SMSs 362, the management switches, 364, and the individual connections 367 illustrated in Figure 17. The individual CSSPs can also communicate directly with each other, as represented by the link 457. One server blade 43-A0 is shown associated with the CSSP 71-A and
5 one server blade 43-B0 is shown associated with the CSSP 71-B. in Figure 20.

For ease of illustration, other components of the system have been omitted, including components in the path between the administrator system 450 and the CSSPs 71-A and 71-B, and other components within the shelf, such as the backup
10 CSSPs, further blades, etc.

It will be noted that the administrator system 450 is illustrated to contain a pool of host identities 452 which can be allocated to the CSSPs 71-A and 71-B, as well as storage for a record 454 of host identities. It also contains storage 454 for
15 recording host identity allocations and transfers.

Figure 20 assumes that the system is in normal operation with a respective set of host identities allocated to the individual CSSPs 71-A and 71-B, for further allocation by those CSSPs to the individual components of the shelf. As illustrated in
20 Figure 20, it is assumed that the CSSPs 71-A has allocated to itself the host identity CNX (420X), and the CSSP 71-B has allocated to itself the host identity CNY (420Y). The CSSP 71-A has a serial number SN0 (460A) and has associated with itself a signature SIG0 (458A). The CSSP 71-B has the serial number SN1 (460B) and has associated with itself a signature SIG1(458B). Similarly, the administrator system can
25 have a serial number SN2 (460Z) and an associated signature SIG2 (458Z). The serial numbers are unique to the respective CSSPs 71-A and 71-B, and are allocated at manufacturing time. The serial number for the administrator can also be allocated in an appropriate manner, and will be unique. The signatures can be allocated by certification authority, and can be used by the administrator system and the CSSPs to

sign messages to enable secure communication of management information. The source and destination host systems 71-A and 71-B can be operable to record transfer operations in respective records 462A and 462B, respectively.

- 5 In the first example, it is assumed that the administration system requires a host identity BNX0 that is allocated to a blade associated with the CSSP 71-A that has developed a fault to be re-allocated to a blade associated with the CSSP 71-B. As indicated above, however, other possible reasons for wishing to reorganize the distribution of the host identities could be to take advantage of different models of
- 10 blade having higher or lower performance and/or components of a different configuration.

A first method of transferring the host identity in a secure and reliable manner will now be described with reference to Figure 21.

15

Figure 21 illustrates an example of an operation to securely and reliably transfer a host identity from a first host system to a second host system.

- In step 510, the administrator system 450 establishes a secure connection to a
- 20 first host system 71-A. In this example, the first host system is the CSSP 71-A, which then operates as a source host system. The secure connection can be established in any conventional manner, for example by a secure socket layer (SSL) connection, or the like, over the connection represented schematically as 456.

- 25 In step 512, the administrator system 450 passes one of the host identities from the pool of host identities 452 as an additional host identity to the first host system (CSSP 71-A) also with a command to use the additional host identity to enable the transfer of the host identity BNX0 (hereinafter referred to in this example as a first host identity) from the blade to 43 to which it is allocated to the blade 43-B0.

In step 514, the first host system 71-A recognizes from the table 440A that it is the blade 43-A0 that has allocated to it the host identity BNX0 and establishes whether that blade is still operational. The first host system 71-A is operable to check whether the blade 43-A0 is still active. If it is, then the first host system 71-A can issue a command to the blade 43-A0 to shut down its operating system and to hold that blade in a reset state, thereby to relinquish the host identity BNX0. However, as such a host identity transfer is typically performed in response to the failure of a blade, the blade will typically already have shut down. As a result, the host identity will typically already have been relinquished to the pool of host identities 440A held by the source host system 71-A. If the blade 43-A0 is not faulty, or can be restarted, the first host system is then able to use the additional host identity to reconfigure the blade 43-A0, and to restart the blade 43-A0 using the additional host identity. Alternatively, if the blade 43-A0 cannot be restarted for any reason, it is maintain in a reset or quiesced state until it is replaced and/or repaired, at which time the additional host identity is available to configure the blade concerned.

Although in the example operation of step 514, the operating system is shut down, as an alternative in some examples, the blade could be quiesced, for example by shutting down the network connections, instead of or as well as shutting down the operating system for the replacement of the host identity. Quiescing the blade rather than shutting it down can prevent invalid use of the host identities while still enabling a quick restart.

Accordingly, it can be seen that in step 514, the first host system effectively reconfigures itself to use the additional host identity.

In step 516, the first host identity (BNX0) is passed to the administrator system via the secure link. This can be done by the administrator system reading and

5 caching the first host identity, and then the first host system 71-A deleting the first host identity from its pool of host identities and 440A and passing this to the administrator system. Where the administrator system is operable to read and cache the first host identity before it is passed from the first host system to the administrator system, the risk of a loss of a host identity due to a failed transfer can be avoided. The administrator system can buffer the received first host identity in the buffer 462.

10 In step 518, the administrator system terminates the secure connection to the first host system and establishes a secure connection to the second host system 71-B. As before, the secure connection can be established in any conventional manner.

15 In step 520, the administrator system passes the first host identity (BNX0) from the buffer 462 to the second host system 71-B via the secure link. This can be done by the administrator system in two stages, with a copy being sent and cached by the second host system before transferring the first host identity. As before, the risk of a loss of a host identity due to a failed transfer can thereby be avoided. The first host identity is transferred with a command to use the first host identity BNX0 for the blade 43-B0.

20 In step 522, the second host system is then operable to use the first host identity as the host identity for the blade 43-B0. The first host identity can be added to the pool, or table 440B of host identities in the second host system. The second host system 71-B can recognize from the table 440B whether the blade 43-B0 is currently operational. If it is operational, for example under the host identity BNY0 shown in Figure 20, then the second host system 71-B can issue a command to the blade 43-B0 to shut down its operating system and to hold that blade in a reset state. This will enable to host identity BNY0 to be relinquished and returned to the pool of host identities 440B in the second host system. The second host system is then able to use the first host identity BNX0 to reconfigure and restart the blade 43-B0.

As with step 514, as an alternative to shutting down the operating system, the blade could be quiesced, for example by shutting down the network connections, instead of or as well as shutting down the operating system for the replacement of the host identity. Quiescing the blade rather than shutting it down can prevent invalid use
5 of the host identities while still enabling a quick restart.

Accordingly, it can be seen that in step 514, the second host system effectively reconfigures itself to use the additional host identity.

10 Optionally, in step 524, the second host system could be operable to pass the second host identity (CNY) to the administrator system 450 if this is needed. Alternatively, the second host system could hold the second host identity in its own pool 440B until it is requested by the administrator system.

15 In step 526, the administrator system terminates the secure connection to the second host system 71-B.

If the second host identity is returned to the administrator system in step 524, then in step 526, the administrator system can also be operable to place the returned
20 second host identity CNY into its pool 452 of available host identities for subsequent use.

Alternatively, as represented in Figure 22, the administrator system 450 can be operable in step 526 to buffer the second host identity BNY0 in the buffer 462. Then
25 the steps illustrated in Figure 22 can be performed. For example, the first host identity BNX0 could have been passed to the second host system with a command not only to use the first host identity BNX0 for the blade 43-B0, but also to return the host identity BNY0 (to be transferred to the blade 43A0).

In step 530, the administrator system 450 once more establishes a secure connection to a first host system 71-A. As before, the secure connection can be established in any conventional manner.

- 5 In step 532, the administrator system 450 passes the second host identity (BNY0) from the buffer 462 to the first host system 71-A via the secure link. This can be done by the administrator system in two stages, with a copy being sent and cached by the second host system before transferring the first host identity. As before, the risk of a loss of a host identity due to a failed transfer can thereby be avoided.
- 10 The second host identity can be transferred with a command to use the second host identity BNY0 for the blade 43-A0.

- In step 534, the second host system is then operable to use the first host identity as the host identity for the blade 43-A0. The second host identity can be
- 15 added to the pool, or table 440A of host identities in the first host system. The second host system 71-B can recognize from the table 440A whether the blade 43-A0 is currently operational. If it is operational, for example under the additional host identity transferred in step 512, then the first host system 71-A can issue a command to the blade 43-A0 to shut down its operating system and to hold that blade in a reset
- 20 state. This will enable to host identity BNX0 to be relinquished and returned to the poll of host identities 440A in the second host system. The second host system is then able to use the first host identity BNY0 to reconfigure and restart the blade 43-A0.

- As with step 514, as an alternative to shutting down the operating system, the
- 25 blade could be quiesced, for example by shutting down the network connections, instead of or as well as shutting down the operating system for the replacement of the host identity. Quiescing the blade rather than shutting it down can prevent invalid use of the host identities while still enabling a quick restart.

Accordingly, it can be seen that in step 534, the second host system effectively reconfigures itself to use the additional host identity.

5 In step 536, the first host system 71-A can optionally pass the additional host identity back to the administrator system, where it can once more be placed in the pool 452 of available host identities for subsequent use. Alternatively, the second host system could hold the second host identity in its own pool 440B until it is requested by the administrator system.

10 In step 538, the administrator system can then terminate the secure connection to the first host system.

During the processes illustrated in Figures 21 and 22, the administrator system 450 can monitor the transfer of the host identities and can retry the transfer in the 15 event that the process hangs, and or can issue an alert. Also, the administrator system 450 is able to record the allocation and transfer of the host identities in the record 454. Similarly, the source and destination host systems 71-A and 71-B can record the transfer operations in respective records 462A and 462B, respectively.

20 In the manner as described in Figure 21, it will be appreciated that a secure mechanism has been described whereby a first host identity can be transferred from a first host system 71-A to a second host system 71-B.

By the further steps shown in Figure 22, it is possible for the host identity that 25 was previously held in the second host system 71-B to be transferred to the first host system 71-A. Accordingly, through the combination of the steps of Figure 21 and 22, it is possible for an exchange of the host identities to be performed.

- As indicated above, the administrator 450 can be part of the management system for the overall multi-processor system illustrated in Figure 17, or alternatively it could be a separate administrator system, which is responsible for distributing licenses to the individual systems. Also, although examples have been described in
- 5 which the host identities for a blade associated with a first CSSP to be passed to a second blade associated with another CSSP, it will be appreciated that the same approach can be taken to transferring the host identities for CSSPs.

Figure 23 illustrates another example of the allocation of a host identity.

10

- In step 550, an administrator system can instruct a CSSP to allocate a spare host identity to a given server blade. This could occur, for example, where a new server blade is added to a shelf. A spare host identity could be supplied by the administrator system with an instruction to a CSSP to allocate the spare host identity.
- 15 As indicated above, it is possible for a CSSP to hold spare host identities for allocation to a server blade. Accordingly, where spare host identities are already held by a CSSP, an administrator system could alternatively instruct a CSSP to allocate a spare host identity that it already holds to a given server blade.
- 20 In step 552, the CSSP could be operable to check that a host identity has not already been allocated to the blade concerned.

- If the server blade concerned is not active, then, in step 554, the master CSSP could allocate the spare host identity to the server blade concerned and then start that
- 25 server blade using the spare host identity.

If a host identity had already been allocated to the server blade and/or the server blade was already active, then as part of step 554, the master CSSP could optionally be operable to query the instruction received from the administrator system.

Otherwise, or if the instruction were confirmed in step 554, then in step 556 the master CSSP could be further operable to issue a shutdown command to the server blade to cause it to shut down and thereby to relinquish the host identity it was using.

5 The master CSSP could then allocate the spare host identity to the server blade concerned in step 554 and then start that server blade using the spare host identity.

As an alternative to shutting down the operating system in step 554, the blade could be quiesced, for example by shutting down the network connections, instead of

10 or as well as shutting down the operating system for the replacement of the host identity. Quiescing the blade rather than shutting it down can prevent invalid use of the host identities while still enabling a quick restart.

The host identity relinquished by the server blade could then be held by the

15 CSSP as a new spare host identity, or until return of the relinquished host identity is requested by the administrator system. Alternatively, the host identity relinquished by the server blade could be returned directly to the administrator system.

Where copies of the host identities are stored in separate components, for

20 example in the back-up CSSP and/or in the front indicator board, then, in the methods of Figures 21 to 23, the master CSSP can be responsible for updating those copies by rewriting the copies of the host identities accordingly.

In the examples described above, it is indicated that information is passed

25 between the host systems and the administrator system using secure socket layer transactions. However, it is not necessary to use secure socket layer communications. In other examples of the invention, it may be chosen to effect all communications as encrypted messages encoded using a selected encryption algorithm and a predetermined parameter or parameters.

For example the administrator system could be operable to read a property of the source and destination host systems (e.g., the serial number SN0 and SN1 from the first and second CSSPs 71-A and 71-B) as parameters to be used for encoding the messages. The appropriate parameters could then be given to hosts concerned. When the administrator system or a host transmits a host identity, it could be operable to encode the host identity using, for example, the serial number of a recipient system and an appropriate encryption algorithm. The selection of a particular encryption algorithm could be a matter of choice according the desired degree of security. For example a public/private key encryption algorithm could be used. A message including the encoded host identity to be sent could further be encoded using parity codes, error correcting codes (ECC) or cyclic redundancy checks (CRC), for example for verifying correct transmission. When the administrator system or a host receives one or more messages including a host identity, it can then be operable to decode the received message, for example using its own serial number and a decryption algorithm appropriate to the encryption algorithm used. Where appropriate, any signature used can be verified as well. The recipient host system could also be operable to check the validity of the decoding operation, for example using any parity, ECC or CRC used.

It may not be necessary to resort to the use of signed messages or secure socket layer transactions if a network in which the invention is implemented is reasonably secure. Through the selection of an appropriate encryption algorithm, a desired degree of security could be provided as will be apparent to one skilled in the art. Enhanced security, from multiple security measures (e.g. encryption plus signed messages and/or secure socket layer transactions) should only be needed where there is concern about malicious user spoofing of systems. Similarly, the use of parity, ECCs or CRCs or indeed any other approach to message integrity checking is a matter of routine choice according to the requirements of any particular implementation.

Thus, there has been described a system where a first host system can be configured initially with a first host identity and a second host system can be configured initially with a second host identity. A host identity can belong to only one host system at a time. An administrator system can hold an additional host identity.

- 5 The administrator system can securely connect to the first host system and can pass the additional host identity to the first host system. The first host system can reconfiguring itself to use the additional host identity and can pass the first host identity to the administrator system. The administrator system can securely connect to the second host system and can pass the first host identity to the second host system.
- 10 The second host system can then reconfigure itself to use the first host identity and pass the second host identity to the administrator system. In this manner the first host identity can be passed securely from the first to the host system without a risk of duplication.

- 15 It will be appreciated that the dynamic transferring of host identities in the manner identified above enables dynamic configuring of a blade system with the advantage that the configuration of the blade system and the allocation of blade system resources to multiple users can be facilitated. Thus, it is possible, for example, to operate a virtual local area network (VLAN) in a dynamic and readily
- 20 reconfigurable manner over a blade system. Portions of the resources of the blade system can be allocated to different users, and then the allocation of resources can be modified dynamically to the system users according to changing system requirements.

- The particularly described example is merely illustrative of the present
- 25 invention, and is not intended to be limiting. It will also be appreciated that many modifications are possible within the spirit and scope of the present invention. For example, although in the present instance the first and second copies 420 and 430 of the host identities 410 are stored in the CSSPs 71, this need not be the case, and they could instead be inserted in other components within a multiprocessor system.

However, in the described example, where the CSSPs 71, and in particular the service processor 74 within the CSSPs 71, has the role of providing management functions within the shelf, it is advantageous for the copies of the host identities to be associated directly with the service processors 74 of the CSSP 71, and for the service processor
5 to perform the steps described above for verifying the integrity of the host identities.

Also, although in the described example, a complete copy of the set of host identities is held in the service processors that form support units, in other examples one subset or respective subsets of the host identities could be stored in the respective
10 support units according to a desired implementation:

In the examples described above, each server blade includes a single processor and is configured as a single host system. However, each server blade could include multiple processors. The processors could be configured as a single host system, or
15 alternatively they could be configured to provide multiple host systems.

Indeed, it will be appreciated that other embodiments could have a very different structure. For example, although examples have been described based on server computers, the host systems could be formed by respective client systems,
20 interconnected via a network, for example via the Internet.

Accordingly, although the embodiments above have been described in considerable detail, numerous variations and modifications will become apparent to those skilled in the art once the above disclosure is fully appreciated. It is intended
25 that the following claims be interpreted to embrace all such variations and modifications.